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

Participants

Alan Mackey	Director, IGT
Sir Christopher Kelly	Chair, Responsible Gambling Strategy Board
Chris Percy	Lead Researcher, BetBuddy
Claudio Corradini	Managing Director, Accenture
Clive Hawkswood	CEO, Remote Gambling Association
Dirk Hansen	CEO, GamCare
Graham Weir	Worldwide Head of Responsible Gambling, Ladbrokes
Dr. Gregory Slabaugh	Senior Lecturer (Associate Professor), City University London
Jon Watkin	Manager, Responsible Gambling Strategy Board Secretariat
Lynda Atkinson	Head of Social Responsibility, Genting Casinos UK
Maris Bonello	Integrity Analytics Manager, Unibet
Paul Hope	Programme Director for Consumer Policy, UK Gambling Commission
Roger Parkes	Global Head of Compliance, Betway
Simo Dragicevic	CEO, BetBuddy
Dr. Tillman Weyde	Senior Lecturer (Associate Professor), City University London

Introduction

The roundtable was convened to enable industry representatives to debate Responsible Gambling (RG) algorithms, which are becoming increasingly important in regulatory compliance and managing customer experience. The roundtable was structured around four key discussion themes:

1. **Model Accuracy v. Transparency and the Need for Human Oversight**
2. **Future Applications of Artificial Intelligence and Analytics**
3. **Modelling and Analytics in Financial Services: Lessons Learned**
4. **Frameworks for Model Development and Testing.**

Introducing the roundtable, Simo Dragicevic (BetBuddy) contextualised RG algorithms into the broader context of the UK's RG strategy, outlining four key pillars that underpin the strategy:

- **Harm Prevention:** Broad awareness strategies aimed at the population
- **Harm Minimization:** More specific programmes and tools aimed primarily at those at higher risk of experiencing harm
- **Problem Gambling Treatment:** Referral and treatment for those who are suffering harm
- **Research:** Building an evidence base to inform efforts to improve harm prevention, harm minimization, and problem gambling treatment.

The primary focus of the roundtable was the application of analytics to support Harm Minimization strategies in gambling. The event coincided with the publication of two peer-reviewed publications by BetBuddy and the Research Centre for Machine Learning at City University London (City) in 2016 in International Gambling Studies (Vol. 16, Issue 2) and the 22nd European Conference on Artificial Intelligence.

We shouldn't underestimate the importance of understanding early pathways to harm in relation to the gambling population lower down the pyramid. It might be that there was some learning from analytics which could be valuable for these players too

Sir Christopher Kelly
Chair, RGSB

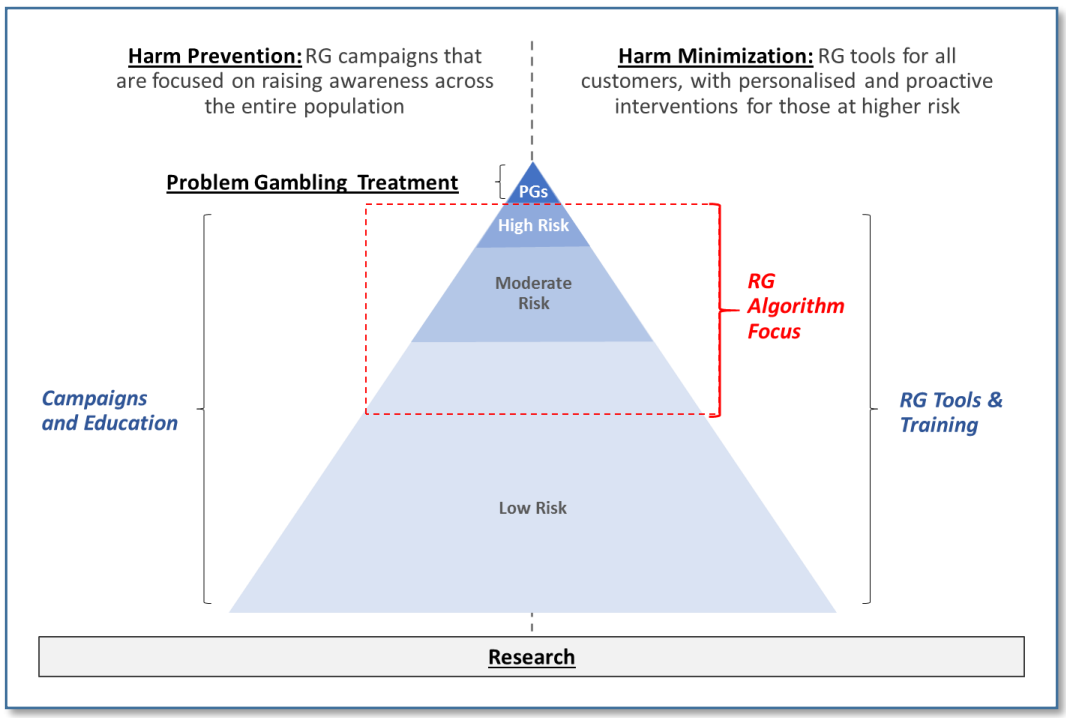

Harm Minimization as 'Below The Line' Marketing

The difference between Harm Prevention and Harm Minimization approaches were likened to Marketing 'Above the Line' (ABL) and 'Below the Line' (BTL) strategies. ABL is where marketing is undertaken at the macro level, and include conventional media as we know it, television and radio advertising, print as well as internet (i.e. Harm Prevention approach). BTL is marketing at the micro level, is more one to one, and can involve highly targeted direct marketing focused on conversations (i.e. Harm Minimization approach).

The major benefit of RG algorithms is the ability to use data to personalise the RG experience and tailor the gambling experience and interventions to those who need them most.

This is particularly pertinent as practical experience shows problem gamblers (PGs) do not perceive themselves as PGs, meaning generic RG awareness campaigns and current RG tools may have limited impact on this group.

Whilst there was broad agreement in the role that RG algorithms can play to deliver a more effective Harm Minimization strategy, Christopher Kelly (RGSB) noted that we shouldn't underestimate the importance of understanding early pathways to harm in relation to the gambling population lower down the pyramid. It might be that there was some learning from analytics which could be valuable for these players too.

Instinctively, greater model understanding would be a higher priority compared with greater accuracy
Paul Hope
Programme Director, UK Gambling Commission

1. Model Accuracy v. Transparency and the Need for Human Oversight

The topic of algorithm accuracy v. transparency was presented by Chris Percy (BetBuddy). There are domains where algorithm accuracy will always trump the need for any transparency, for example in the fields of image recognition, where most businesses do not care how the algorithms work, but they do care that they work well. Likewise, an eCommerce retailer would probably never need to explain to a customer why a product recommendation was made to them.

The need for accuracy in RG algorithms however is considered a second priority compared to the need for understanding. 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. When this question was posed to the roundtable participants, instinctively the consensus was for a more understandable algorithm over a more accurate one.

Dirk Hansen (GamCare) stressed the value of interpretability, especially as this can be an advantage when providing treatment as the counsellor has specific and relevant behavioural indicators to discuss. Alan Mackey (IGT) stated that a 100% accurate algorithm would be preferable, but because this is highly unlikely in this context, even a slightly less than 100% accurate algorithm leaves room for doubt, therefore focusing solely on accuracy may not be the right approach.

From a regulatory perspective, Paul Hope (Gambling Commission) stated that whilst the regulator is mostly focused on the outcome rather than the process, without having looked at this particular point in detail, instinctively, greater model understanding would be a higher priority compared with greater accuracy. Roger Parkes (Betway) added that transparency was important to enable an operator to manage any regulatory challenge with regards to a model that had been deployed.

Interpretability can be an advantage when providing treatment as the counsellor has specific and relevant behavioural indicators to discuss
Dirk Hansen
CEO, GamCare

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

Chris Percy (BetBuddy) described results from BetBuddy's and City's latest research to be presented at the European Conference on Artificial Intelligence (September, 2016). The research successfully applied TREPAN, a rule extractor, to a trained neural network that could predict self-exclusion. The results converted a 500+ parameter neural network model into a 9 route decision tree with only a 1 percentage point loss of model accuracy.

This raises the real possibility of implementing RG algorithms that offer both high accuracy and high transparency and interpretability, with a traditional machine learning model being used for prediction, and a second TREPAN-like model being used to explain and interpret results. Roundtable participants agreed that it would be very interesting if there was the opportunity to apply a two tier approach, using two algorithms that provided both the accuracy and interpretability.

The challenge to gather model 'ground truth' or event data

We converted a 500+ parameter neural network into a 9 route decision tree with only a 1 percentage point loss of model accuracy. This raises the real possibility of implementing RG algorithms that offer both high accuracy and high transparency and interpretability
Chris Percy
Lead Researcher, BetBuddy

The issue of the lack of 'event data' for effective modelling (i.e. problem gamblers and their associated data) was discussed in depth, what Dr. Tillman Weyde (City) referred to as 'ground truth' data. Self-exclusion is widely used across industry for RG model development because it is easier to collect this data compared with say, gambling screen data for all customers. Whilst in the UK Retail gambling sector self-exclusion is widely regarded as a good proxy for harm, in the internet sector, self-exclusion does not always translate to a customer experiencing harm, as the self-exclusion process can be a convenient and near instant way for a customer to close their account if they do not wish to gamble on a site for a variety of reasons. In contrast in the retail gambling sector, the customer has to undertake a more detailed, considered, and thus time-consuming process to self-exclude.

Roger Parkes (Betway) highlighted that because self-exclusion is not a perfect proxy for harm, getting the right model input data is very difficult for operators. This also points towards looking at other support options, such as excellent customer service or intuitive site design, to make up for limitations in data analysis.

Maris Bonello (Unibet) said that researchers continue to face the major of challenge of not being able to look at player behaviour across multiple operators. It was noted that many internet self-exclusions come from players with very short tenures after registration, which Graham Weir (Ladbrokes) confirmed was an issue facing other internet operators.

Clive Hawkswood (RGA) said that increasing industry's understanding of model input variables and the key harm triggers was a key objective of the research being undertaken by PWC for the Responsible Gambling Trust (RGT). Whilst this will accelerate industry's understanding of an appropriate set of model inputs, which will resolve some operator modelling challenges, the research will not combine customer play across different internet operators due to privacy hurdles. This was felt by some to be a missed opportunity to tackle understanding customer behaviour across multiple operators, a research question yet to be tackled.

Whilst a 100% accurate algorithm would be preferable, because this is highly unlikely in this context, even a slightly less than 100% accurate algorithm leaves room for doubt, therefore focusing solely on accuracy may not be the right approach
Alan Mackey
Director, IGT

Looking beyond self-exclusion in internet gambling

We have opportunities to look beyond self-exclusion data, such as how customers interact and use a broader variety of RG tools, which can help deepen understanding of harm
Graham Weir
Worldwide Head of Responsible Gambling, Ladbrokes

In a discussion to find alternatives to self-exclusion as the main proxy for harm in internet gambling, Maris Bonello (Unibet) explained efforts at Unibet to start to analyse account closures with follow-up telephone calls to customers to understand their reasons for account closure. Graham Weir (Ladbrokes) highlighted the opportunities to look beyond self-exclusion data, such as how customers interact and use a broader variety of RG tools, which can help deepen understanding of harm. The idea of giving customers a shorter screen to take (i.e. less onerous than the 9 point PGSI questionnaire, for example), during the customer sign-up process online, or at other early customer interaction opportunities, was proposed as an alternative means to collect event data.

Simo Dragicevic (BetBuddy) suggested that the brief screens, such as the Brief Biosocial Gambling Screen (BBGS), developed by the Division on Addictions (a Harvard Medical School affiliate), could help. The BBGS, a three-point screen, had produced good results in testing (e.g. sensitivity of .96), and therefore may be more palatable to customers compared with the longer PGSI.

However, the question of when to ask a customer to take such a screen was discussed as collecting such data has traditionally posed challenges e.g. if taking a PGSI-type gambling screen was mandatory (e.g. on sign-up) there was a risk that many customers would complete the process as quickly as possible, rather than taking their time to answer the questions candidly.

This led to a further discussion on the challenges of how to effectively use the outputs of models. Roger Parkes (Betway) stressed that the modelling of harm, despite the known challenges, is an easier problem to solve compared with what's the best way to use the model insights to make customers change their behaviour, a discussion point that was re-visited later.

2. Future Applications of Artificial Intelligence and Analytics

An overview of future trends and developments in AI was presented by Dr. Tillman Weyde (City) and covered the following themes:

- **Multimodal learning:** Combining different forms of data: transactional, text, video, audio, and sensors
- **Emotion recognition:** Examining the use of deep learning algorithms for emotion detection
- **Intelligent agents:** The development of AI that will manage more human tasks independently
- **Connected devices:** The Internet of Things, and the integration and connections of devices (wearables, home appliances, and industrial goods and machines)
- **Data security, privacy, and ethics:** A look at the challenges that come when we generate more and more data in everyday life.

If loss of control data sets were available then researchers may find better behavioural indicators than those currently known
Sir Christopher Kelly
Chair, RGSB

Understanding loss of control using AI

A key discussion was centred around the used of AI to detect emotions, and particularly whether loss of control could be detected during game sessions. The ability to model simple emotions is already providing positive results and is now available, such as through Microsoft's Emotion API. However, the key challenge to making emotion recognition work is getting access to a sufficiently large number of training data sets. Modelling emotions such as 'happiness' and 'sadness' can be considered relatively easy to detect given the abundance of training data, however getting lots of training data sets that exhibit 'loss of control' could prove to be challenging. Chris Percy (BetBuddy) discussed how the case of Google's automated image tagging algorithm, which automatically tagged a black person as a Gorilla, highlighted some shortcomings of AI and the dangers of relying solely on it.

Manufacturers are testing the use of cameras in gaming machines and the initial focus is primarily based on improving the customer experience [rather than RG]
Alan Mackey
Director, IGT

Maris Bonello (Unibet) suggested that as well as a future algorithm that could model 'loss of control', recognising more simple emotions such as 'anger' or 'disgust', particularly if used in a multimodel capacity in conjunction with play data such as when a customer is losing, could be helpful in signalling problem gambling.

If in the future loss of control could be detected accurately, then this might provide opportunities to make more direct interventions
Sir Christopher Kelly
Chair, RGSB

Simo Dragicevic (BetBuddy) said that a technical challenge with real-time analysis of multimodal data is the high dimensionality of data that is assessed. Also it was noted that emotions can change fast, vary from individual to individual, and can also be controlled.

There is a significant difference, both in implications and the potential interventions, between someone who is feeling anger but has controlled it and is continuing rationally, and someone who is unable to control their feelings of anger.

Christopher Kelly (RGSB) stated that if loss of control data sets were available then researchers may find better behavioural indicators than those currently known. If in the future loss of control could be detected accurately, then this might provide opportunities to make more direct interventions, such as an automatic timeout.

Alan Mackey (IGT) said that he was aware that machine manufacturers such as IGT have been testing the use of cameras in gaming machines. The initial focus was primarily based on improving the customer experience, however he wasn't aware of uses cases looking at gambling harm yet. Lynda Atkinson (Genting) said as there are higher levels of staff on a casino floor they are better placed to monitor behaviour, applying customer oversight as part of their customer service duties, so at this stage in its conception, AI is less likely to give the benefits it might in other sectors.

Using AI to model cultural differences

The discussion evolved into applying AI to understand cultural differences in gambling behaviours, and whether or not the feasibility to understand these existed today. Dr. Gregory Slabaugh (City) said that whilst this is possible today, the data requirements to enable this to be done effectively will be high. For example, it would be possible to build a self-exclusion prediction model for a subset of UK Retail gamblers who are from a distinct culture (e.g. Far East Asian problem gamblers living in the South East of England), however the gathering of sufficient model training sets that provide the necessary 'ground truth' would be needed. It was also noted that the process of collecting and labelling data is time consuming (and thus expensive).

Industry has the challenge to decide whether they build multiple RG models for different customer sub-types, or to try to standardise on a single model that is applicable to different types of gamblers
Simo Dragicevic
CEO, BetBuddy

Simo Dragicevic (BetBuddy) outlined that industry has the challenge to decide whether they wish to build multiple RG models for different customer sub-types, or to try to standardise on a single model that is applicable to different types of gamblers.

This is a cost/accuracy trade-off, and whilst there are instances where separate models will almost certainly be needed if a high level of accuracy is required, such as if the game domain and characteristics are very different, it may be appropriate to apply a general model in some circumstances.

The BetBuddy and City's self-exclusion prediction model, published in International Gambling Studies, produced a model accuracy of 87% using a random forest technique. This model was developed as a general model using 5 broad, although highly configurable, risk indicators that had applicability to multiple domains. The value of this approach is that operators can apply such general models to more than one game type or jurisdiction, which keeps model development and maintenance costs lower. Clive Hawkswood (RGA) said that such an approach may be more attractive to operators who are struggling to remain competitive in the light of increasing compliance costs.

Spotlight on Research

INTERNATIONAL GAMBLING STUDIES, 2016
http://dx.doi.org/10.1080/14459795.2016.1151913

Routledge
Taylor & Francis Group

Predicting online gambling self-exclusion: an analysis of the performance of supervised machine learning models

Christian Percy[a], Manoel França[b], Simo Dragičević[b] and Artur d'Avila Garcez[b]

[a]BetBuddy, London, UK; [b]Department of Computer Science, City University London, London, UK

ABSTRACT

As gambling operators become increasingly sophisticated in their analysis of individual gambling behaviour, this study evaluates the potential for using machine learning techniques to identify individuals who used self-exclusion tools out of a sample of 845 online gamblers, based on analysing trends in their gambling behaviour. Being able to identify other gamblers whose behaviour is similar to those who decided to use self-exclusion tools could, for instance, be used to share responsible gaming messages or other information that aids self-aware gambling and reduces the risk of adverse outcomes. However, operators need to understand how accurate models can be and which techniques work well. The purpose of the article is to identify the most accurate technique out of four highly diverse techniques and to discuss how to deal analytically and practically with a rare event like self-exclusion, which was used by fewer than 1% of gamblers in our data-set. We conclude that balanced training data-sets are necessary for creating effective models and that, on our data-set, the most effective method is the random forest technique which achieves an accuracy improvement of 35 percentage points versus baseline estimates.

ARTICLE HISTORY
Received 2 July 2015
Accepted 31 January 2016

KEYWORDS
Responsible gambling; gambling self-exclusion; machine learning; supervised learning algorithms; problem gambling; Internet gambling; three-tier model

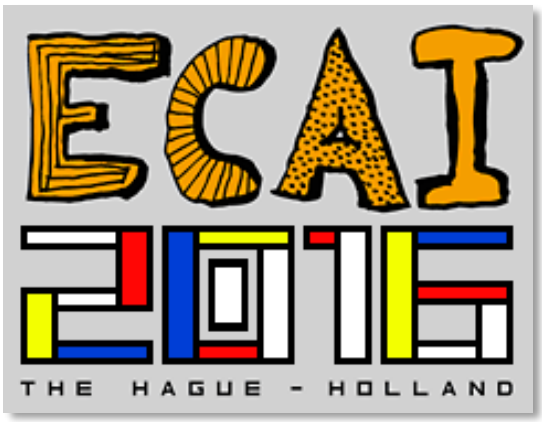

Since 2014 BetBuddy has been collaborating with academics from the Research Centre for Machine Learning at City University London, to build predictive models and to externally benchmark and publish BetBuddy's RG models in an effort to further research and continuous improvements in this domain.

Two research papers have been published in 2016. The first paper, published in International Gambling Studies (Vol. 16, Issue 2), benchmarked BetBuddy's general self-exclusion prediction model against open research data sets from Harvard's Division on Addictions, producing model accuracies of 87% using the random forest technique.

The second paper, accepted for publication and presentation at the 22nd European Conference on Artificial Intelligence, successfully applies knowledge extraction techniques to make blackbox neural network algorithms readable and interpretable.

Applying intelligent agents to real-time play analysis

In the context of intelligent agents, Simo Dragicevic (BetBuddy) explained how predictive models could be developed to analyse real-time data to make customer interventions e.g. i) the analysis of potentially harmful anonymous play, such as loss of control on Fixed Odds Betting Terminals (FOBTs), and ii) in the context of changing the game dynamics in real-time, such as a scenario where a customer who plays very intensively on a high stakes game triggers an AI agent to subtly slow-down the speed of the game.

The general feeling from the roundtable participants was that the application of intelligent agents to FOBTs had potential given that most play on FOBTs (c.85%) is anonymous, making traditional analytics redundant. The ability of an intelligent agent to continuously analyse and detect patterns of play, whilst feasible, will require the collection of training data for model development, which will be a key challenge.

Such an approach was considered worth pursuing given other options (e.g. lower stakes, mandatory player cards, mandatory loss limits, etc) are considered unattractive to industry for a variety of reasons. Paul Hope (Gambling Commission), when asked, couldn't see a credible case for relaxing stake sizes if such intelligent agents existed on FOBTs in the current environment, even if they showed potential in minimizing harm.

Clive Hawkswood stated that whilst the idea of using intelligent agents to change the customer experience could be useful, it is important to strike a balance. Making games unattractive can have the effect of driving customers to other channels or sites where player protection is not as advanced. Maris Bonello (Unibet) said that Unibet has begun analysing the effects of different bonus types on game play, such as whether certain bonuses produce certain behaviours that could be a pathway to harm.

We analyse different bonus types on game play, such as whether certain bonuses produce certain behaviours that could be a pathway to harm
Maris Bonello
Integrity Analytics Manager, Unibet

Whilst the idea of using intelligent agents to change the game experience could be useful, it is important to strike a balance. Making games unattractive can have the effect of driving customers to other channels or sites where player protection is not as advanced
Clive Hawkswood
CEO, RGA

Is RG analytics endangering duty of care?

The increased spotlight on RG algorithms raised the issue of whether their use warranted a greater duty of care on the operators. The roundtable did not discuss any technical legal definitions or applications of duty of care, rather potential broader implications of customer oversight from the use of RG algorithms.

Operators now have more customer information available to them through the use of RG algorithms, however telling a customer that they have been 'detected' was not considered to be the right approach. This is particularly pertinent given the earlier discussions on model accuracy v. model interpretability.

For example, a customer could ask why were they not stopped, or why their accounts were not closed earlier, if they were flagged either chasing losses or chasing winnings.

As analytics improve, it was stressed that operators need to be careful of how it is presented to users, to ensure that an unreasonable and unfulfillable new duty of care expectation is not created. As yet, this concern is not holding back innovation, but it should be monitored at an industry level. Paul Hope (Gambling Commission) stated he was encouraged to see that a duty of care risk wasn't holding back innovation in RG algorithms.

We are encouraged to see that a duty of care risk isn't holding back innovation in RG algorithms
Paul Hope
Programme Director, UK Gambling Commission

3. Modelling and Analytics in Financial Services: Lessons Learned

A key focus of the roundtable was to examine how the Financial Services industry has developed analytics and models for decades. Claudio Corradini (Accenture and Financial Services Executive), who has previously held senior finance positions, explained how analytics and models are extensively used within the financial services industry, covering domains from risk management and customer behavioural scoring, to capital management and algorithmic trading.

Claudio Corradini, Managing Director with Accenture's Financial Services practice, has previously held a number of senior finance positions, including CEO or Barclays Europe (Italy, Portugal, and Spain), CFO of Antonveneta, and CFO of RBS' £300bn Asset Protection Scheme

Company culture trumps model sophistication every time
Claudio Corradini
Managing Director, Accenture

Over reliance on models

Financial services firms had traditionally been given a lot of flexibility to develop internal, complex models to manage capital and risk, primarily to enable the larger and more complex organisations to optimise capital management i.e. free-up more capital to increase trading and lending. However, this played a key part in the credit crisis in 2008 and Claudio Corradini (Accenture and Financial Services Executive) provided a cautionary warning against encouraging a culture of over-reliance on models to make decisions.

As well as the credit crisis, a culture of over-reliance on models in the past to make automated decisions had resulted in banking creating distance between the bank and the customer. For example, when a decision on whether to grant a mortgage to a customer is taken by an algorithm, arguably the needs of customer cannot be fully assessed. Despite this, good process design can overcome such issues, enabling digital formats to offer opportunities for more personalised and closer experiences, even though algorithms are employed.

In financial services, it is typical to implement a 'three-line defence' development and governance process, effectively triangulated by 1) internal compliance or finance teams, 2) internal audit teams, and 3) by external experts
Claudio Corradini
Managing Director, Accenture

Management oversight of complex models

Claudio Corradini (Accenture and Financial Services Executive) stressed the point that company culture trumps model sophistication every time. There existed a risk of 'moral hazard' within organisations that develop and use analytical models. Where firms have minimal oversight of their algorithm development, they may fall foul of an institutional (even if unconscious) bias towards flagging up lower than optimal levels of risk, since it allows greater lending in finance or greater customer marketing in gambling. While these may prop up profits short-term, they are not in the long-term interests of the industry. Moral hazard was a key contributor to bankers taking positions and risk that were not in the best long-term interests of customers and shareholders.

In financial services, it is typical to implement a 'three-line defence' development and governance process to enable effective model oversight to safeguard against such risks.

Oversight of model development and implementation is effectively triangulated by 1) internal compliance or finance teams, 2) internal audit teams, and 3) by external experts. Clive Hawkswood (RGA) said that there is a precedent in the gambling industry for external testing and certification, such as with RNG game testing and also with the GamCare accreditation for RG policies and processes. Therefore, a similar process the RG model oversight and validation could be appropriate in the future.

Claudio Corradini (Accenture and Financial Services Executive) summarised that whilst internal and external processes and validation are significantly important in enabling effective model development and oversight, instilling a culture from the very top of an organisation of putting customer needs first was the single most important factor that would ensure RG models are developed and deployed correctly.

The use of sophisticated v. blunt Instruments

The notion that a complex or more sophisticated model was better than a simple model was challenged by Claudio Corradini (Accenture and Financial Services Executive), who outlined a trend of banking regulators beginning to impose more general, simple models on the industry in a response to reduce risk in the sector. Giving banks 100% flexibility on model development was considered a risky option, and it was debated whether there was a need for regulators to provide more definitive guidelines to industry in the context of RG algorithm development.

The logical place is to start was with more 'blunt' instruments, as provided in the LCCP, but also to build more sophisticated models over time as evidence and analysis produces other reliable predictors. Compared with financial services, the gambling industry is younger in terms of developing sophisticated models of behavioural risk
Roger Parkes
Global Head of Compliance, Betway

The Gambling Commission's Licensing Conditions and Codes of Practice (LCCP) requires the "*provision to identify at risk customers who may not be displaying obvious signs of, or overt behaviour associated with, problem gambling: this should be by reference to indicators such as time or money spent*". Whilst Paul Hope (Gambling Commission) could not give any further direction on the particular point of simple v. sophisticated models, he said he had sympathy for a personalised, player-focused approach, but would want to see the value of this demonstrated empirically.

Roger Parkes (Betway) said that the logical place was to start with more 'blunt' instruments, as provided in the LCCP, but also to build more sophisticated models over time as evidence and analysis produces other reliable predictors. Compared with financial services, the gambling industry was younger in terms of developing sophisticated models of behavioural risk.

It was felt by Maris Bonello (Unibet) that more definitive guidelines on how to implement RG algorithms would certainly help the industry, particularly medium and smaller sized operators who have much less resources compared with larger operators. However, a number of participants felt that moving to a more prescriptive regulatory approach could create and foster a 'box ticking' compliance culture, rather than focusing on developing and investing in better RG algorithms and processes though continuous improvements.

Unibet, who have invested significantly in developing their internal RG algorithms and processes, initially started with a simple model based on communication indicators, and then evolved this over time. Unibet would likely continue to develop and use their proprietary internal algorithm in parallel with any general models or approaches required by regulators.

In a further discussion on establishing general guidelines, some participants felt that it was reasonable to use peer-reviewed research, such as that from the British Gambling Prevalence Survey (BGPS), as a benchmark for establishing risk thresholds for gambling operators who use RG models. There was general support for such an approach, especially given that if operators were reporting numbers to the regulator, discrepancies could be flagged between operators more easily.

We have invested significantly in developing internal RG algorithms and processes, initially starting with a simple model based on communication indicators, which then evolved over time. We would likely continue to develop and use our proprietary internal algorithm in parallel with any general models or approaches required by regulators
Maris Bonello
Integrity Analytics Manager, Unibet

RG analytics as competitive advantage

We look to use our RG capabilities and services as a differentiator when we compete in less well established markets, to differentiate from incumbent competitors
Graham Weir
Worldwide Head of Responsible Gambling, Ladbrokes

The roundtable debated whether better customer RG interactions could lead to a competitive advantage. Claudio Corradini (Accenture and Financial Services Executive) emphasised how trust has been eroded in financial services, and how deploying analytics to understand the customer needs is now the industry's primary focus. Whilst RG is not currently widely viewed as a competitive advantage, in theory it could be if the industry was confronted with a high profile scandal.

As well as trust, one very poor customer experience can be amplified to large audiences via social media, therefore the emphasis on putting the customer first is very much a competitive advantage in retail banking. Banking today is moving away from the mind-set of individual transactions to customer journeys and 'life events', with the aim of understanding and helping customers through their life journeys (as opposed to selling a product).

The casino environment offers greater oversight and arguably greater levels of customer care, and this is something that is of importance to our customers
Lynda Atkinson
Head of Social Responsibility, Genting UK

Lynda Atkinson (Genting) explained how the casino environment offers greater oversight and arguably greater levels of customer care, and this is something that is of importance to Genting's customers. Maris Bonello (Unibet) explained that the RG team at Unibet is focused on implementing RG excellence rather than growing revenues, although there was a large consensus from participants that RG excellence could have a positive impact directly on the brand and indirectly on (increasing) revenues.

Graham Weir (Ladbrokes) explained that Ladbrokes look to use its RG capabilities and services as a differentiator when it competes in less well established markets, to differentiate from incumbent competitors.

In financial services, most models that are developed and used for risk management are highly proprietary and play a key role in helping banks to increase income, therefore it would be inconceivable for competitors to share models and best practice. While the gambling industry is not yet falling victim to these same problems, partly as the algorithms are far more nascent than in finance and not used for revenue generation, it is important to monitor developments and correct for any adverse moves.

It would be interesting to see whether in future, if RG algorithms do play a greater part in customer retention and building brand trust and value, whether operators would be as willing to collaborate on RG initiatives as much as they do today. When asked, operators said they wouldn't 'open source' (i.e. make freely available) their proprietary RG models and algorithms.

Banking today is moving away from the mind-set of individual transactions to customer journeys and 'life events', with the aim of understanding and helping customers through their life journeys (as opposed to selling a product)
Claudio Corradini
Managing Director, Accenture

Can RG be personalised based on gambling experience and ability?

A further discussion centred around the MiFID regulations in finance, which stipulate that if you're an 'expert investor', you are allowed to be sold more complex financial products. Is there a parallel in gambling? There are some indirect parallels, such as being in a more deliberate gambling environment (such as casinos v. pubs) unlocks riskier products, or using a loyalty card in Retail gambling enables fewer restrictions on maximum stakes.

Could these parallels continue and develop a scenario where more astute and experienced gamblers can gamble with less RG restrictions? It was felt that the highly sensitive nature of RG meant that this wouldn't be a consideration for regulators in the short term.

4. Frameworks for Model Development and Testing

The final theme centred on assessing whether a framework could be developed to help the industry to better develop and continuously improve RG algorithms. Simo Dragicevic (BetBuddy) presented a summary of 10 questions which can be used to guide RG algorithm development and oversight (with number 11 added as a further suggestion by the roundtable participants):

RG Model Framework: Development and Assessment Questions	
1	What is the basis and evidence for deciding the algorithmic objective?
2	How was model development undertaken?
3	Are there any potential biases and limitations in the training data sets?
4	How have model thresholds been selected?
5	Has the algorithm been tuned to benefit one set of stakeholders v. another?
6	Is there information that is differentially overemphasized or excluded by the algorithm?
7	Are there any technical limitations in the analytical methods used?
8	What are limits to measuring and operationalizing the criteria used by the algorithm and its outputs?
9	To what extent is the model visible, understandable, and configurable?
10	How often is the model(s) evaluated and updated?
11	What impact is the model(s) having on the objective of minimizing harm?

Operators should be able to use the outputs of RG algorithms to drive various interventions
Paul Hope
Programme Director, UK Gambling Commission

Most of the points summarised in the RG Framework were debated in the earlier sessions, such as: (1) the basis for selecting model development event data e.g. self-exclusion, (2) simplistic v. sophisticated model techniques and their implications, (3) generic models v. specific models, which can take into account cultural factors, for example, (4) using prevalence study research to provide model benchmarks, (5) mitigating 'moral hazard' risk in model development, (6) the current LCCP guidelines on customer behaviours to assess (with emphasis on time and money spent), (7) using 'blunt' v. more sophisticated methods, (9) trade-offs between model accuracy v. transparency, and (10) the case study of developing the right cultures and model oversight frameworks within financial services.

In the field of medical imaging analysis, the concept of 'open data' for research, involving the sharing of data across many research organisations, has had a significantly positive impact on improving disease detection rates
Dr. Gregory Slabaugh
City University London

Testing outcomes is as important as testing model accuracy

The participants were all in agreement in stating that despite the challenges to developing effective RG algorithms, knowing how best to use the outputs (8) and measure the RG models effectiveness in minimizing harm (11) are also significant challenge facing the industry. It was discussed that the nature of the challenge (i.e. are RG algorithms accurate and do they change behaviours) makes empirical analysis of the effectiveness of RG models feasible, and ultimately the effectiveness of a model testable.

The roundtable widely acknowledged the importance of empirically testing not just the predictions of analytics but also the interventions on which they are based and the extent to which analytics can provide information that improves the outcomes of interventions. For example, what changed as a result of using it? How did customer behaviour alter? How are interventions improved as a result? Does it generate insights and outputs that usefully inform investigative reports or clinical support?

Early impact assessment results from Unibet suggest that RG algorithms can have a positive impact on encouraging players to use RG tools (e.g. limits) and in moderating excessive play.

Paul Hope (UK Gambling Commission) stressed the importance of (9) in ensuring that operators can use the outputs of RG algorithms to drive these various interventions.

Are false positive flagging concerns valid?

There is always a cost associated with both false positives and false negatives, and the industry needs to decide on which is the bigger priority
Dr. Tillman Weyde
City University London

Further discussion centred around (5), particularly whether it was justifiable to develop algorithms that benefit the vast majority of customers i.e. non-problem gamblers, rather than focusing on those most at-risk. There was broad agreement across the roundtable that the primary objective of RG algorithms should be to minimize harm. From a model calibration perspective, this would practically mean that placing a greater emphasis on higher model sensitivity (i.e. lower false negatives) was more important than higher specificity (i.e. lower false positives). Dr. Tillman Weyde (City), explained there is always a cost associated with both false positives and false negatives, and the industry needs to decide on which is the bigger priority.

At this point there was no industry evidence to suggest that proactive RG interventions were having a detrimental impact on non-problem gamblers, and any previous concerns may have been over-stated in their seriousness.

It was also noted that when banks make security checks, customers are generally not put-off by these and they have become expected as part of normal business. Banks are just doing their jobs to protect customers.

Therefore, industry should test empirically whether such interventions impact the customer experience sufficiently seriously to alter their playing behaviour or willingness to continue to enjoy and spend on gambling.

Whether RG interventions should be standardised across industry was debated. Clive Hawkswood said that in order to 'normalise' proactive RG interventions the industry could use consistent or standardised approaches (e.g. in player messaging), which could protect against degrading the player experience (a risk should operators take widely differing approaches). It was noted that a potential downside to this could be that RG interactions could become 'too normalised', thus reducing their impact. There was agreement that it was sensible for operators to approach interventions in a constructive and educational manner.

Open research data drives better results

Open research data sets available from the Division on Addictions [a Harvard Medical School affiliate] had proved extremely useful in benchmarking RG model results to a different IGT data set
Simo Dragicevic
CEO, BetBuddy

The final discussion centred around the role of the industry and regulator in making available data sets to help operators to build effective RG models. Dr. Tillman Weyde (City) asked whether it would be possible for the regulator to provide anonymised data examples of harm to aid in model development. Dr. Gregory Slabaugh (City) explained that in the field of medical imaging analysis, the concept of 'open data' for research, involving the sharing of data across many research organisations, has had a significantly positive impact on improving disease detection rates, for example.

Simo Dragicevic (BetBuddy) asked whether industry data could be made more widely available for RG research e.g. RGT research data sets, explaining that the open research data sets available from the Division on Addictions (a Harvard Medical School affiliate) had proved extremely useful in benchmarking RG model results to a different IGT data set. Everyone agreed that this made sense, and whilst commercial sensitivities surrounding the data was the only reason offered why they couldn't be made more widely available, it was noted that strict definitions on data management and usage can resolve this issue.

Christopher Kelly (RGSB) said that the RGSB's research strategy, which was in the process of being updated, was open to new ideas.

Review

Key messages for operators emerging from the roundtable:

1. *Responsible Gambling algorithms can support multiple strategies – not only harm minimization (aimed at moderate and higher risk populations) but also possibilities may exist in harm prevention (aimed at the wider low risk population)*
2. *Start simple or with a single model – build complexity in as and when the company expertise grows and when the business demands greater sophistication*
3. *Look beyond self-exclusion and traditional self-test results (for remote in particular) as the only source of event data to model harm – this requires new processes to capture customer feedback and potentially new data sources e.g. loss of control*
4. *Place emphasis on transparency and interpretability and be cautious about algorithms managing every RG interaction autonomously – AI can still make mistakes*
5. *Company culture beats models always – senior management set the tone as to how customers are treated and this is the best safeguard against 'moral hazard' e.g. increasing marketing activity even if this doesn't benefit the customer*
6. *Triangulate model development and governance – make use of internal teams and external expertise to ensure robust model development and oversight*
7. *Focus sufficient efforts on operationalization – making effective use of the algorithm outputs e.g. such as for player messaging, will require sustained efforts across multiple departments*
8. *Empirically test Responsible Gambling algorithms – test model accuracy and also outcomes e.g. whether customer behaviours are changing following an intervention*
9. *Use research, such as population prevalence studies, to benchmark customer risk scores – this also enables operator-to-operator comparisons of results*
10. *Don't let perceived increased duty of care concerns inhibit innovation in algorithm development and adoption – if implemented carefully, Responsible Gambling algorithms could strengthen the operator's position*
11. *Responsible Gambling analytics might become a greater source of competitive advantage in the future – whilst the primary aim is harm minimization, better insights and interactions brings operators closer to their customers.*

BetBuddy

BetBuddy contacts:

Simo Dragicevic
Telephone: +44 7977 414127
simo@bet-buddy.com

Chris Percy
Telephone: +44 7733 063475
chris@bet-buddy.com

CITY UNIVERSITY LONDON

Research Centre for Machine Learning at City University London contacts:

Dr. Artur Garcez
Telephone: +44 20 7040 5060
a.garcez@city.ac.uk

Dr. Tillman Weyde
Telephone: +44 20 7040 5060
t.e.weyde@city.ac.uk

Dr. Gregory Slabaugh
Telephone: +44 20 7040 5060
gregory.slabaugh.1@city.ac.uk