



Game Play:

The next frontier
for Responsible
Gambling

Three emerging Responsible Gambling dynamics

Responsible gambling (RG) technology is evolving at pace. In just a few short years, industry practices have progressed from SQL databases and Excel to machine learning algorithms. This should be taken as evidence of the seriousness of the industry's commitment to using data and new technology to make gambling safer.

Yet there is huge demand for further innovation, both from within and outside the industry. The Spotify generation have increasingly high expectations for a safe entertainment environment, while the UK Gambling Commission are pushing for continuous improvements. The pandemic has only intensified the call for further action.

In October 2019, the Gambling Commission passed [far-reaching new RG guidance](#) and also issued a number of [challenges](#) to the industry. With the pandemic leading to increased financial strain, isolation and even boredom, concerns for at-risk gamblers led to the Commission providing [updated guidance](#) in May of this year.

These regulations all but mandate the adoption of advanced machine learning to analyse game data and identify customers at risk of harm. Three key themes emerge across these provisions that demand Game Data Science.

1. Markers of Harm Monitoring

Firstly, operators need to improve their monitoring of **Markers of Harm**. Larger operators need “comprehensive systems” that “should draw on all available sources of data to give a comprehensive picture of the customer’s gambling”. All operators must use “a range of indicators based on research, experience and shared practice” that should include “play indicators” and “time and spend indicators”. Of the six indicators that the Commission explicitly recommend that operators track, only one pertains to self-exclusion and the use of gambling management tools.

2. Real Time

Secondly, larger operators should be **continually monitoring** their customer base and making *timely interventions*. In the October 2019 Commission guidance, “in-play real time monitoring to identify harmful behaviour as it occurs” is discussed as an optional measure that operators might take to protect customers. Citing research indicating that high-risk customers tend to gamble overnight, the Commission emphasised that “*unmonitored overnight gambling carries an increased risk*”. The May 2020 Commission guidance emphasises that operators need to “ensure the continual monitoring of their customer base — identifying patterns of play, spend or behaviours [that] have changed in recent weeks”.

At the time of writing, the Commission are undergoing a consultation into tougher rules for identifying and tackling gambling harm. The Commission have already stated that, as part of this consultation, they will be looking to operators to flag risk more expeditiously.

3. Safer Game Design

A third emerging theme pertains to **safer game design**. The Betting and Gaming Council in the UK should be applauded for having taken the initiative with a **Game Design Code of Conduct** that includes significant commitments to minimum game cycle speeds, and the removal of turbo play, slam stops and multi-slot play. But it is unanimously agreed that research and understanding around game design is continually developing, and we do not yet have a complete understanding of which features may be more likely to lead to potential risk, or conducive to safer play.

These three fundamental areas — comprehensive tracking of behavioural Markers of Harm through gameplay data, timely or real-time intervention and safer game design — will form part of the next chapter in the story of RG technology evolution. That chapter is already being written today.

The state of play

The overwhelming majority of people who gamble do so safely. For these players gambling is a social activity, a challenge, a way to relax; their play is entertaining and sustainable.

However, a minority experience harm. The latest data from **the NHS Digital Health Survey** indicates that the **prevalence of problem gambling is 0.5%** (approximately 280,000 people in England). A further **0.8% of gamblers are at moderate-risk**, and **2.7% are classified as low-risk**. The figures for the UK appear to be in line with **Europe-wide estimates**, with the rates of **disordered gambling ranging from 0.1-0.8%**, with an additional **0.1-2.2% demonstrating potentially problematic play**.

Rates of problem gambling vary widely between different gaming activities, with participation in slots having one of the **highest proportions of problem gambling at 9.2%** according to the Gambling Commission.

Covid-19 has, of course, significantly impacted consumer behaviour and the gambling industry at large. **Preliminary data** released by the UK Gambling Commission after the introduction of government-mandated social distancing indicated that while the majority of gamblers did not increase the time or money spent on gambling, the most engaged gamblers did in fact increase their gambling activity.

The volume of activity on slots in particular grew **25% year-on-year in March 2020**. More **recent data** from August shows that the number of online slots sessions lasting longer than an hour — an RG metric that the Commission are focused on — decreased by 7%. Anthem's data indicates that just **2.5% of slots sessions now** last longer than one hour. Over the same period the number of customer interactions **increased by 11%**. This can be taken as tentative evidence of the success of the industry in proactively preventing harm during a volatile and challenging period.

Risk bands are not static: **longitudinal research** demonstrates that **9-15% of individuals** in moderate- or low-risk categories will transition to problem gambling within a few years. **Research** conducted in Australia indicates that **70+% of gambling harm** is experienced by people whose behaviour is not classified as problem gambling. These data points suggest that timely detection and mitigation of moderate- and low-risk gamblers can have a significant impact on harm reduction.



The background of the page is a photograph of the Golden Gate Bridge in San Francisco, California. The bridge's iconic orange-red towers and suspension cables are visible against a clear blue sky with a few wispy white clouds. The bridge spans across the top and right side of the page, with a large, semi-circular teal shape overlapping the bottom and left sides, creating a modern, graphic design.

Why responsible gambling?

The industry is tackling problem gambling and gambling-related harm with seriousness, energy and creativity. And with good reason: the ethical and economic incentives are clearly in place.

Fines issued by the UK Gambling Commission for failure to protect problem gamblers have increased over time, **with over £30m in fines levied and 11 licenses revoked in the 2019-20 period.**

Beyond reducing the risk of fines and regulatory actions, there are also positive incentives for achieving responsible gambling excellence. There are customer satisfaction and loyalty benefits: after all, gambling is largely enjoyable and sustainable. There are also intangible but important reputational benefits: safer gambling is not a nice-to-have corporate social responsibility (CSR) item; today's consumer has increasingly high expectations that companies will be trustworthy and ethical, and they will reward operators that meet those expectations. Identifying vulnerable players also allows operators to market more freely to the rest of their player base.

RG is about more than just detecting and intervening on problem gambling — it is also about designing games and curating gaming environments to promote safety. When an engineer builds a bridge, safety is not an afterthought or optional bonus feature: it's integral to the design. Similarly, the gambling industry is taking huge strides in making safety increasingly integral to game design. In the near future, the “safe” in “safe game design” may sound as superfluous as it does in “safe bridge building”.

Reducing the number of players playing in a problematic fashion by encouraging safe betting behaviour and usage of responsible gambling tools isn't just important in an ethical sense. Players who bet in a safe affordable manner are more valuable long term for studios and operators.

Three dimensions of innovation

The incentives for continued RG innovation are clear, but the challenges that still lie ahead are considerable. We now have the capability — breakthroughs in cloud technology and software systems and the migration of AI from academia to enterprises — to produce more powerful RG solutions that can meet the challenges to support the industry to the next level of RG practice in line with regulations and to support players.

There are three fundamental areas in which the technological vanguard is catching up with or pulling ahead of the pace of regulatory demand.

From 2D to 3D player profiling

Current approaches often focus on predicting customers who are liable to self-exclude. Whilst prior self-exclusion data is tangible and useful, it is only part of the solution. Regulators are now mandating that operators also track behaviours that are indicative of harm. Future Anthem have combined respected cognitive science research and state-of-the-art machine learning technology to build models that can detect a more three-dimensional problematic player profile. We can now apply these models to surface subtle behavioural Markers of Harm hidden in gameplay data.

From delayed to real-time deployment

Updated guidance from the Commission requests that operators “continuously monitor” customers, while empirical analysis indicates that high-risk players are disposed to bet out-of-hours during the week. Analysing player data in batches and making delayed interventions is no longer best practice. Future Anthem’s tech stack can plug into operators’ ecosystems, ingesting millions of spins in real time and mitigating player risk as it arises.

Responsible game design

Regulators have issued a challenge to the industry to embed safer gambling principles into games. The industry has in turn taken the initiative to make a number of major commitments to responsible game design. But it remains an open question whether mathematical game properties and game features may lead to Markers of Harm, and if so to what extent. Future Anthem are developing answers to these questions, by fusing markers-of-harm models with granular game intelligence.

From Excel to Machine Learning

Technological progress in this space has been salutary, however meeting the current and emerging regulatory demands for risk detection is highly challenging.

In a short space of time, the industry has progressed from manual-intensive detection to AI systems in production.

“When I first started in the gambling industry our responsible gambling model existed solely on an antiquated laptop, only looked at a handful of datapoints and was run once a week. Here at Future Anthem we employ clusters of super computers to analyse each bet in real-time.” Dr Edmond Mitchell, Lead Data Scientist at Future Anthem

The prevailing approach is to use machine learning algorithms that are trained to identify gamblers who voluntarily self-exclude from a platform. These algorithms are valuable: they can flag at-risk users who would benefit from self-exclusion before they feel the need to do so.

However, models engineered to detect self-exclusion face limitations. Many problem gamblers never end up self-excluding, often due to addiction or lack of awareness. Equally, some gamblers who don't have a gambling problem self-exclude.

Indeed, a **prominent analysis** of 10,000 UK customers found that **80% of self-defined problem gamblers have never used a self-exclusion tool, and only 31% of those that have self-excluded in the past self-define as a problem gambler.** Because self-exclusion is a very noisy indicator of problem gambling, models that have been trained to detect self-exclusion may generate large numbers of false positives and negatives.

These machine learning systems that are currently widely used in industry represent important progress beyond Excel and SQL databases. But if operators are to meet the emerging regulatory demands introduced above, these models need to be augmented with algorithms that are tuned in to the subtle behaviours that are truly diagnostic of problem gambling.

Beyond self-exclusion: 3D player profiling

Operators will increasingly need to build technology that can identify behavioural markers of gambling-related harm. But which markers should be used?

A broad range of analyses converge on the following four fundamental categories:

- **Frequency**, which relates to the number of sessions a player engages in during a specified period.
- **Intensity**, which relates to the duration of a player's session and the volume of betting behaviour within that session.
- **Variability**, which relates to how a player's gambling behaviour varies over time.
- **Trajectory**, which relates to whether these three key markers are increasing or decreasing over time which is key to identifying change in behaviour.

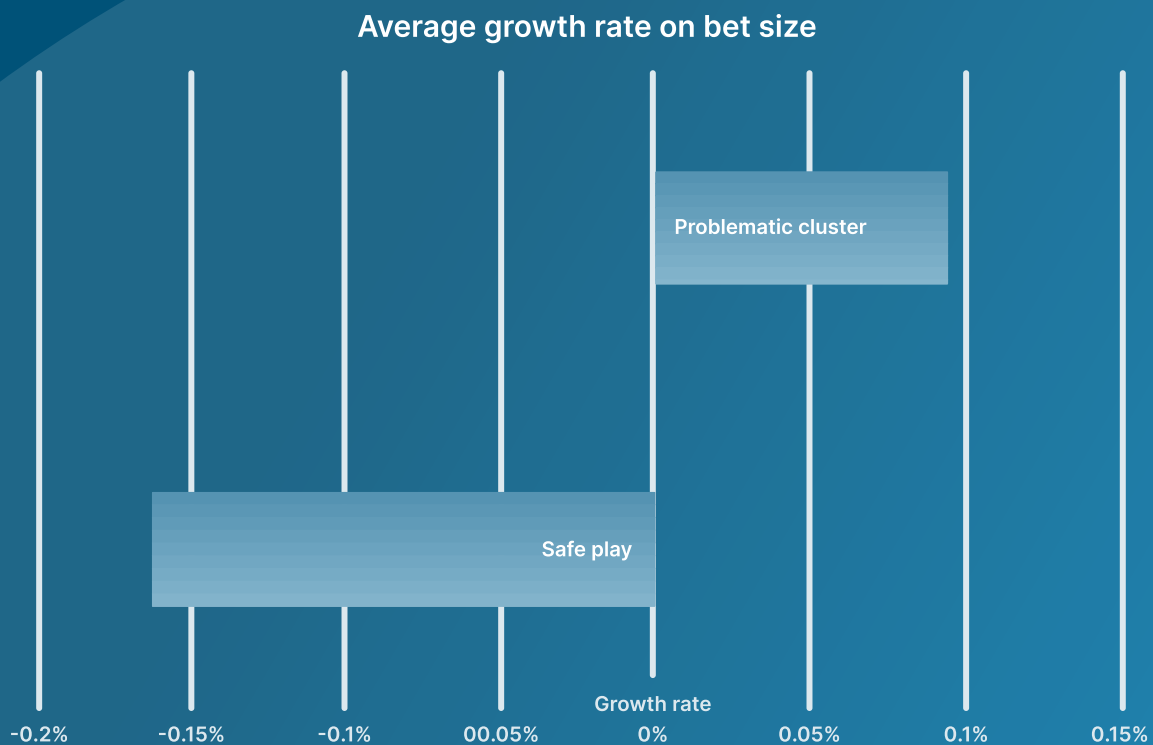
Behavioural patterns that are known — from academic research and clinical studies — to be diagnostic of problem gambling fall under these four categories. For example, simultaneously playing many games is a “frequency” phenomenon, playing extremely long sessions is an “intensity” marker, loss-chasing/erratic-play falls within “variability”, while deviating from regular playstyle would be detected under “trajectory”.

The challenge facing tech makers, a challenge raised by the needs of the customer and issued by the regulators, is to build models that can detect these key behavioural Markers of Harm. From a data science standpoint, this is in part an advanced feature engineering problem.

At Future Anthem, we have placed the four fundamental behavioural categories — Frequency, Intensity, Variability, Trajectory — at the heart of our risk identification models. Access to spin-by-spin game data allows us to engineer highly discriminative and novel features that are designed to capture the full sweep of validated Markers of Harm.

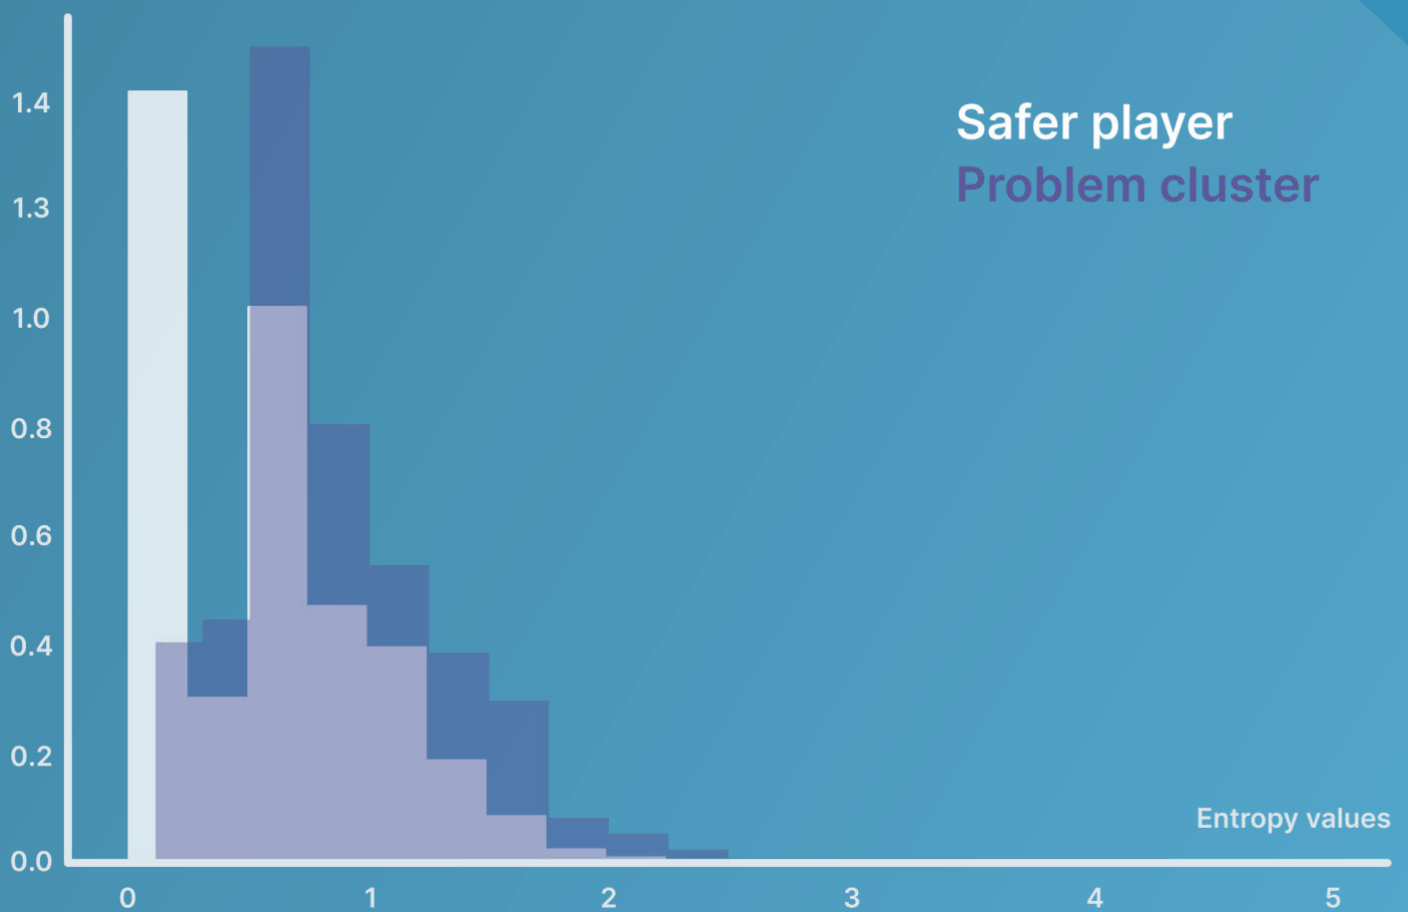
For example, let's look a little closer at one of these behavioural patterns: loss-chasing, which is the tendency of a gambler to increase betting in an attempt to recoup losses. This phenomenon is widely cited as a **defining characteristic** of problem gambling, with one **prominent theoretical model**, the "pathways" model, positing that loss-chasing is common to all subtypes of problem gambler.

Utilising information extracted from spin level data allows the growth of stake levels to be measured during a gaming session. This feature combined with other novel features allows Machine Learning Models to accurately identify when a player is loss chasing.



The image above shows the average growth rate for two clusters Future Anthem has created. The safe players are much more likely to have a negative stake growth rate. That is to say that their bet size is more likely to decrease over time. In contrast the problematic cluster is more likely to have a positive growth rate in stakes over the course of their session. This metric on its own does not tell the full story of what is happening in a game session however in combination with other variability metrics it is quite useful.

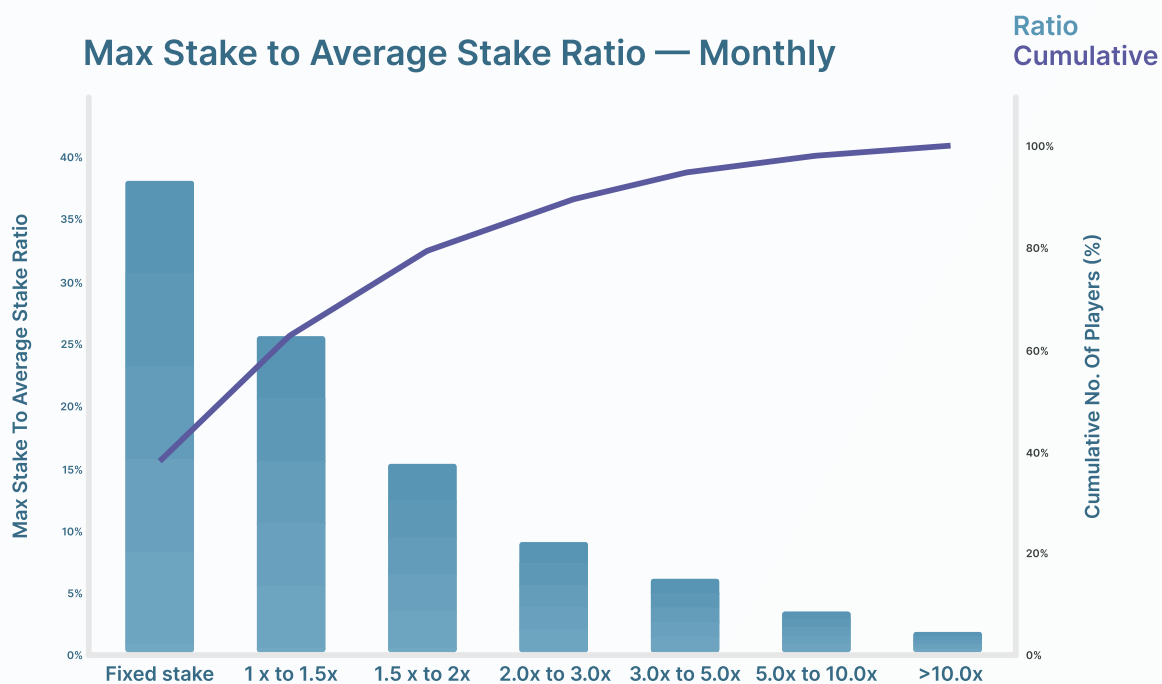
To give a second example, we have found that entropy, a measure of randomness of a series of bets, gives insight into play volatility. For example, a player who consistently stakes at the same level would have zero entropy, whereas a player whose stakes fluctuate wildly would have high entropy. We've found that this metric helps clustering algorithms identify variable play and thus help detect problematic gambling behaviours (see graph below).



Another key metric that we extract from the data is how we compare each bet with the player's average stake. In the graph below you can see that almost half of all players do not change their stake and **over 80% of players** do not wager more than twice their average in a **30-day window**. **Less than 6% of players will wager more than five times their average**. This is not to say that everyone who has **bet more than 5 times their average** has a problem but more that it helps paint a picture of a player's behaviour. Knowing when a player is playing outside their normal behaviour is hugely helpful to the machine learning algorithms that detect Markers of Harm.

These are just a handful of examples of the features that need to be extracted to allow machine learning clustering algorithms to be able to isolate and identify Markers of Harm. For instance, Future Anthem currently uses over 80 features and the list is continuing to grow. Usage of state-of-the-art machine learning clustering algorithms allow these features to be combined in multiple dimensions. Scoring highly on one metric does not necessarily indicate problematic play, it is when a player scores highly across each of the four Markers of Harm that we are confident the player is playing problematically.

All this critically useful information is hidden in gameplay data. By feeding it into machine learning pipelines, we can create 3D behavioural player profiles of a customer base. We can now use this data to really know our customers and we have the opportunity to intervene before serious harm is done.



From Past Time to Real Time

Problem gambling doesn't sleep, and RG technology shouldn't either.

As we've seen, regulators are pushing for continual monitoring and ever-faster interventions when problems are detected. Anthem's data indicates that problem gamblers are 4 times more likely to move from low to high-risk, meaning that there is an opportunity for operators to get in front of the problem and intervene early. Operators that want to get ahead of the regulations and that aspire to RG excellence are starting to innovate with solutions that can identify harm and intervene *in real time*.


We now have the cloud computing technology and machine learning software to make this happen. Future Anthem's tech stack can:

- Ingest and process massive volumes of gameplay data simultaneously and in real time
- Push that data to predictive models, using a live service and containers, that can detect risk on a spin-by-spin basis
- Use a real-time scoring API to deliver predictions and recommendations back to operators

Problem gambling happens 24-7, and effective RG solutions should be able to protect customers any time and all the time. We have got the technology to do this today.

Here are two portraits of slot players that illustrate the player-traits that the Anthem model learns and interventions it recommends – and how you should be considering your player base.

Our algorithms learn similarly nuanced and fine-grained details about every customer and build highly personalised plans for each.





Player portraits

Using the approach outlined in this paper allows for one of more fundamental requests from the UKGC to be addressed. Previously it was very difficult to detect low income players who were betting outside their norm. Trajectory metrics extracted from the data allow for the quick detection of changes in behavioural patterns. Below we have some examples of players who would not normally be correctly classified in current problem gambling algorithms.

Noa

Game behaviour

- Regularly stakes between £4,000 and £6,000 in a session
- Typical session length is around 2 hours
- Game loyal - plays two games regularly
- Noa's betting patterns are consistent across sessions
- The time that elapses between spins is consistent, as is the time he takes between sessions
- He places a relatively low number of stakes above his median stake size

Risk profile

- High stake play would trigger typical models to falsely classify Noa as high-risk.
- Markers of Harm modelling identifies Noa's low variability and even trajectory play, placing him in a safe cluster.
- Affordability, and source of funds, are of course separate considerations

Eto

Game behaviour


- Shorter number of spins and shorter session time
- Her usual stake size is around £200
- In her next session stake size doubles
- The wait between sessions was much shorter than usual
- There has been a high degree of volatility in betting patterns
- There has been an upward trend in stake size
- She places large amounts of bets over her median

Risk profile

Markers of Harm models detect Eto's volatile play and risky trajectory, flagging her as a player that needs a level of intervention.

Markers of Harm models enable operators to avoid needlessly intervening with Noa's gaming experience.

More importantly, they identify the risk that Eto's gameplay is flagging and enable operators to preemptively intervene with, for instance, on-site pop-up messaging, customer support or account-blocking.



Get ahead of the game

It is widely accepted that there is much work to do in how games are designed from an RG perspective. The Betting and Gaming Council and its studio participants have begun this journey with the industry's first **Game Design Code of Conduct** published in September of this year.

The first phase commits to the removal of game features prone to abuse such as: turbo play, slam stops and the ability to play multiple games simultaneously. The subsequent phase aims to enhance transparency by implementing the ability to easily differentiate between "wins" below the stake size and those actual wins which are larger to their stake size.

Safer game design remains an open research topic. As stated by the Betting and Gaming Council, the game design code is a living, breathing document. There are complex challenges facing the industry: we need to understand not just how individual features contribute to play that exhibits Markers of Harm; we also need to develop statistically robust tests to investigate the ways in which features may combine to lead to Markers of Harm.

With the advent of technology that allows vast amounts of spin-by-spin data to be scrutinised at once, this can be achieved. Game features can be isolated and their impact on a player's ability to safely gamble measured. How does volatility affect a player's ability to play safely? How does the RTP or the number of reels? Does increasing or decreasing minimum or maximum stake have a detrimental effect on a player? Does a combination of certain hit rates and bonus modes encourage safer play? These are questions that state-of-the-art technology platforms which have access to spin by spin data can answer.

It is unanimously agreed that research and understanding around game design is continually developing, and we do not yet have a complete understanding of which features may lead to risk, or are conducive to safer play. At Future Anthem, we are turning our unique Game Data Science perspective and Markers of Harm algorithms to solving these open questions in responsible game design.

Amplify your game with Game Data Science

We have profiled three emerging responsible gambling areas in which there is latent demand for AI solutions and machine learning models that can analyse gameplay data: 3D player profiling to detect risk, real-time risk identification, and safe game design.

In terms of detecting risk in game data, Future Anthem have the solutions. Anthem's tech stack has already processed more than 9 billion transactions from over 2 million unique players. We have built machine learning models that can drill inside gameplay data to surface subtle behavioural Markers of Harm. Our systems can ingest high volumes of data and deploy models in production in real time, seamlessly integrating with an operator's data ecosystem.

Of course, other challenges remain, from affordability checks to cross-operator player profiling, each requiring different data assets.

Game Data Science is a critical and important aspect of the rapidly developing frontier of RG technology innovation. Our goal is to innovate and create with data, taking a progressive approach to improve industry best practice and the playing experience for everyone.

If you'd like to find out how we listen to player and game behaviours to provide solutions on the cutting edge of Responsible Gambling, [get in touch today.](#)



Amplify your game