

Paper 379-2011

Patron Analytics in the Casino and Gaming Industry: How the House Always Wins

Scott Sutton, The Venetian and Palazzo Resort Hotel & Casino, Las Vegas, NV

ABSTRACT

To stay ahead in the ultra-competitive Las Vegas Strip, casinos have to engage in predictive analytics to segment and market to their patron databases. Predicting the future worth of gamblers is the ultimate goal of analytics in the casino industry. The gaming industry is inherently suited to advanced analytics due to the vast amount of data and information collected within the casino industry. This paper will discuss how to leverage your patron database to effectively segment and market to your customers. Specifically, it will discuss how an analysis of patron spending in gaming, hotel, and retail can be used to predict the future value of your patrons. Furthermore, once you've identified those patrons, we'll discuss how to optimize your marketing to those patrons.

INTRODUCTION

Casinos in gaming hubs like Las Vegas, Reno, and Atlantic City face stiff competition from one another for customers' money and loyalty. In 2010 there was over \$8.9 Billion in gaming revenue for Clark County, \$5.7 Billion of which was from the Las Vegas Strip alone (Las Vegas Convention and Visitors Authority, 2010). With more competition and rooms than ever, casinos are putting a heavy emphasis on analytics to compete with one another and ensure they attract the best and most profitable guests. Modern casino analytics developed along with the transition to electronic tracking of gaming transactions. Patron management systems have provided the gaming industry with considerable amounts of highly detailed data about when, where, how often and how much patrons are playing. In the 1980's and 1990's, casino patron loyalty programs, originally called "slot clubs", started popping up in many of the larger casinos. These slot clubs encouraged customers to sign up for player cards and, in return for loyalty to the casino, patrons would receive rewards such as complimentary rooms, access to special events, and other offers. This was revolutionary, as it allowed casinos to track gaming behavior down to the individual level, leading to more accurate information about patrons' gaming behavior and interests. The information could then be used to better segment customers, predict future behavior, and improve marketing outcomes. As casino analytics advanced, casino resorts started incorporating the relevant data from hotel, dining, retail, entertainment, and other outlets to get a more complete view of patron's behaviors. A recent development is that many of the major gaming loyalty programs, especially those in competitive markets such as Las Vegas, are now also rewarding non-gaming spending in order to encourage customers to keep non-gaming spending at their respective properties, in addition to providing additional data about non-gaming behavior.

The vast amount of data that modern patron rating systems and loyalty programs supply analysts can be overwhelming. Intense hardware, time and labor costs are required to dig through the considerable amounts of gaming, hotel, and retail data and find relevant information. Even the smallest casinos will have an abundant amount of transactional data about their patron's gaming behaviors. Data mining tools, such as SAS Enterprise Miner™, can be used to minimize the time and labor costs involved in finding useful information in piles of data and improve the models and metrics used to build marketing campaigns. If you know what to look for and where to look, advanced statistical training is not needed to produce significant results to implement and improve marketing campaigns. All of the analytical methods mentioned in this paper can be executed using SAS® software. For the mass market segment of patrons that don't have personal relationships with hosts, patron analytics are essential for maximizing revenue driven by mass market marketing campaigns. This paper intends to provide some strategies and solutions to two of the most important questions in casino patron analytics – who are the best patrons (customers)? And, what are the best offers for these patrons? More precisely, this paper will focus on providing insight into the ways that casinos can improve marketing strategies by addressing some of the core components of these questions, such as:

- How much is a patron worth, how much can we expect a patron to lose in the future, and who are the most valuable patrons?
- What patrons come together?
- What patrons are most likely to abuse an offer?
- What patrons are most and least likely to respond to an offer?
- What offers perform the best?

HOW MUCH IS A PATRON WORTH?

Most experts in the industry would agree that determining a patron's worth is the first and foremost responsibility of patron analytics in the casino industry. These offers hinge on the accurate prediction of a patron's behavior on that

Patron Analytics in the Casino and Gaming Industry: How the House Always Wins, continued

next trip. Naturally, predicting a patron's future behavior is a very complex task that is affected by a number of variables, many of which are outside factors that the business might not have insight into, including total income, expendable income, ethnicity, reasons for a trip (convention vs. vacation), etc. Although that information is often available to append through third parties, there is still plentiful information found with in-house data that can be used to build models and metrics to predict a patron's future worth. Once the worth of a patron is determined, patrons can then be segmented into groups based on other behaviors and effective marketing campaigns can be developed around those behaviors.

First, it is important to determine what worth is, as the definition of worth is critical for deciding how valuable a patron is and how much to reinvest in the patron in the future. There are two main components of worth – the financial sources of worth (i.e., gaming or hotel) and the unit of time to which it refers (daily, weekly, monthly, etc.). Additionally, worth can refer to historical worth, which is already known, or future worth, which is unknown. The definition of worth will likely depend on both the various financial sources of revenue that affect the business directly and the exact business problems that are being addressed. Gaming worth can also be broken down into various sources (i.e., slots, tables, poker, race and sports) depending on the business issue being addressed. All casino analytics departments should have a solid method for predicting the various types of patron worth based on the sources and time periods they need for making informed marketing decisions. For instance, daily gaming worth would be most useful for building a campaign with a daily free play offer. For example, two separate patrons, each with an annual worth of \$10,000, will have completely different daily worth if one patron comes one day a year and the other comes one day a week. Alternatively, annual worth might be more useful than daily worth to identify the patrons to target for an exclusive event. In other words, aggregate measures of worth are ideal for identifying the most valuable patrons to the business overall, whereas individual measures are usually suited more toward identifying valuable patrons for specific offers and campaigns.

Most revenue sources are fairly straightforward – room revenue is how much the patron paid for a room, restaurant revenue is how much they paid for food and drinks, etc. Gaming revenue, however, is a little more complicated because probability is involved. There are two important measures used to assess a patron's gaming worth - actual and theoretical loss. Actual loss is how much money the patron actually lost (or won), whereas theoretical loss usually refers to the amount of money a patron is expected to lose based on the amount of money wagered, the time spent playing, and the probability associated with type of games played. Theoretical loss tends to be more heavily relied upon for predictive analysis and is a much stronger predictor of future behavior, as actual loss is usually used to measure campaign performance and profitability. Below are the general formulas used to calculate theoretical loss for tables and slots:

- **Table Theoretical Loss = Average Bet x Time Played x Speed of Game x House Advantage**
- **Slot Theoretical Loss = Coin in x Hold Percentage**

PREDICTING FUTURE WORTH

It is important to note that there are a number of issues with the measurement and accuracy of these factors. Some are nearly impossible to control for, such as the fact that patrons must show or use their loyalty card when playing for the data to be tracked to the player. Most table ratings are complicated and prone to error because much of the rating is manual. Although some properties now have table games that incorporate RFID technology in tables and chips to more accurately track individual wagers and time played, this technology is rare. Slot ratings can be complicated in that it is difficult for most machines and rating systems to differentiate between free play and cash play. Before diving in to predictive analysis, it is worth the effort to investigate how your property(ies) handles some of these measurement issues and determine whether there are ways to increase the quality and amount of rated data. However, these measures are normally accepted as reliable and accurate measures of gaming worth.

Once patron worth has been defined, the business can then use data mining and modeling to estimate predicted worth in the future. Simple metrics based on historical behavior, such as Average Daily Theoretical Loss or Average Trip Theoretical Loss, will produce fairly accurate predictions of future worth. However, advanced predictive models are able to predict worth with more accuracy and power by accounting for both patterns in behavior over time and relationships between predictive inputs that exist within casino data. There are a variety of techniques that are used to develop models to predict future worth, the most common being regression models. **Multiple regression** models are the most common because they utilize a variety of predictors and the relationships between those predictors to predict future worth. For example, a model built to predict future gaming trip worth might be generated based on historical information about theoretical win, actual win, credit line, time on device, nights stayed, and average bet. Regression models can also be built using such categorical variables as predictors, as gender, ethnicity, age range, or other demographic variables. Developing separate models based on categorical variables, such as separate models predicting worth for slot and table players, might produce models with less error and better predictions. Regression models are particularly effective because the model can be used to score historical data to predict an unknown outcome, which is worth in this case, within a certain degree of confidence.

IDENTIFYING THE MOST VALUABLE PATRONS

In addition to developing models to predict future worth, there are other analytical methods to determine a patron's value to the business. One way to identify the best patrons is to try and separate the skilled gamblers from the unskilled. Common sense tells us that the unskilled gamblers should be worth more overall. Most casino databases won't have a really good measure of skill, such as how often a blackjack player hits on 16 when the dealer is showing 6 or less. However, it is possible to look at whether a patron is usually a loser or winner. A quick and easy way to evaluate a player's skill is by calculating the percentage of trips where the player actually lost money. For instance, did a player with five trips lose money on all five of those trips? Although this might just be an indicator that the patron will play until he is out of money or time, it also is a fairly simple way to identify the patrons that do not come away as winners very often. This is an instance where actual loss might be a good predictor of worth, as we would rather have these patrons in the casino. This is highly useful for games requiring skill, such as table games and video poker, but might not be very useful for slot players where there is little skill involved and constant losses are more likely to be attributed to luck.

Although slot machines are not really skill based, we can still differentiate between patrons by looking at the strategies and behaviors of slot players. One quick and easy way to separate slot patrons is compare how much play they have on participation machines relative to owned machines. Since casinos have to pay a certain percentage of win or handle to the slot manufacturer for participation games, patrons that primarily play non-participation games are slightly more valuable to the casino. A slightly more complex metric for slot players is to look at their average bet relative to the maximum bet on the games they play. Usually, the maximum bet has to be played in order to be eligible for jackpots and progressives. Given two patrons of similar theoretical worth, the one that plays closer to the maximum allowed bet is more likely to hit a jackpot than the one who doesn't. Usually the patron with the higher average bet would seem to be more valuable, but since the lower bet patron is less likely to hit a jackpot, the lower bet patron might be a lower risk. This metric could be useful on its own, or could be used as either a predictor in a model for future worth or a decision tree predicting whether a patron will respond. These are just a few examples of how data mining, along with predictive modeling, can provide useful information to differentiate between players that might otherwise seem very similar.

IDENTIFYING PATRONS THAT COME TOGETHER

Another important consideration in the discussion of patron worth is household worth. This refers to the combined worth of multiple patrons that tend to make their trips together. This can be difficult to identify, as these patrons might stay in one room or separate rooms, or one patron might only come when accompanied by another patron. Additionally, the other patron might make trips without the first patron. Although identifying household worth can be tricky, it can pay huge dividends by helping to account for revenue that looks like two separate individuals but can be combined into one "household". Many patron management systems contain the functionality to link accounts so that patrons that come together (i.e., married couples) can be easily identified. Unfortunately for the casino analyst, patrons might not be allowed to have linked accounts because of system limitations or because of business policies based on tax and gaming regulations. Nevertheless, data mining can be used to identify groups of patrons that come together without linked accounts. First, we need to identify patrons that make their trips at the same time as one another. Second, we can use a combination of various data points to identify the "households", such as:

- Last name (great for identifying relatives that come together)
- Address (roommates or patrons living together with separate last names)
- The room or floor that the patrons stayed in (patrons that come together tend to request rooms near each other)
- City and State (for friends/relatives from the same area)
- The time of day that games are played (whether they are playing at the same time)
- The type of games played (whether they are playing in the same location or close to one another)
- Restaurant/retail charges (whether they have charges from the same outlet on the same day)

The above method is a great way to identify patrons with trips over the same period of time that have some commonality about their behavior that we can use to be fairly certain the patrons are together. For some of the more subjective measures (i.e., room, floor, city, time/type of play) it's a good idea to be more conservative about how many overlapping trips the patrons have. For instance, two patrons with one overlapping trip and rooms next to each other may or may not be in the same "household" group. However, those same patrons with 5 overlapping trips, each with rooms next to one another, are much more likely to be in the same household grouping. In this manner, household grouping can identify a group of four patrons that are of "middle of the pack" worth individually, but come together and stay in the same room every time and thus are worth more as a group. Now we can adjust our marketing efforts and send a better based on their combined worth and the knowledge that we're not really marketing to four unique individuals, rather to a group of related patrons.

IDENTIFYING PATRONS AT RISK OF ABUSE

Patron Analytics in the Casino and Gaming Industry: How the House Always Wins, continued

Predictive models of worth will likely take into account the factors that predict whether a guest will play on a future trip. It is also a good idea to build a separate model to identify patrons that are likely to use a future offer and not play at all. Since many offers in the casino industry tend to be for complimentary rooms that are given to patrons upfront, patrons that redeem offers and do not play have a considerable impact on campaign success and profitability. Therefore, it is important to identify the patrons that are likely to do that and to adjust their offers based on that knowledge. **Decision trees** and **logistic regression** are common statistical methods used to identify patron characteristics that predict the likelihood of a patron (or segment of patrons) to abuse an offer.

Some factors that are likely predictors of abuse are age (younger patrons are more likely to abuse), gender, and history of abuse. Additionally, survey data (e.g., from follow-up surveys after a patron's visit) that is linked to individual patrons can be used to identify other predictors. If a patron thought they were treated unfairly and had a bad experience in the past, they might take an offer for a free room as revenge for that bad experience. When they come in they take all the perks available from the offer then walk across the street to gamble and spend money so they can later get a better offer from across the street. By identifying the patrons at risk of abusing offers, the business can decide how to market to those risky patrons. For instance, someone might reach out to the patron to try and rectify the situation if they had a bad experience. Instead of sending them the general offer for a free room, they would be sent an offer that requires them to play to a certain level or they will have to pay for their room. However, it might be best to not even give them the option to get a complimentary room. In this case, it is helpful to know what offer is the best type of offer to send a patron and whether it's even worth the money to send them the offer at all.

OPTIMIZING OFFERS

In addition to predicting the future worth of patrons, it is important to know which marketing campaigns are the most effective for driving response, revenue, and profit. In general, certain offers are better than others, and specifically certain offers will be better for certain patrons. The common components of marketing involve offers for rooms, restaurants, retail, and gaming (i.e., free play). While knowing the probable future worth of a patron is critical for determining the reinvestment level for which a patron is eligible, patrons' behaviors and interests can be used to identify the offer(s) that will be most appealing to each patron and generate the most profitable response. Offers that include free rooms and gaming free play are historically the strongest drivers of response. However, a free room is not the ideal offer for every patron. On one hand, some patrons won't be eligible for a free room because their predicted gaming worth is too low to warrant a free room. On the other hand, not every patron that's eligible for a free room has to be offered a free room to respond – some might be willing to pay for a discounted room or even a full price room. By analyzing the likelihood that a patron will respond to a certain offer or offers, casino analysts can optimize the offer that each patron is given in order to maximize the amount of revenue and profit driven by marketing campaigns as a whole.

The most basic way to identify the best offer is through **A/B testing**. A/B testing involves testing two different offers against one another in order to identify the offer that drives the highest response and the most revenue/profit. More advanced statistical methods can be used to generate likelihood of response scores and classification scores. Some of the more common statistical approaches are **logistic regression**, **decision trees**, and **discriminant analysis**. Essentially, these statistical methods use historical data to find the factors that are related to whether a patron responds. Those factors can then be used to assess the likelihood of response based on the similarity of a patron profile to that of responders. These methods have historically been used in direct marketing analysis to identify the best types of offers and the most likely responders. In order to build accurate and predictive response models, historical data about response is required. The likelihood of response might be a broad measure of response that refers to the likelihood a patron will respond to any offer, or it might be specific to the likelihood of response to a specific type of offer. Additionally, it's a good idea to select test segments of patrons for the purpose of continually testing new offers. Doing so will help to ensure that there is a large amount of response data that can be used to build models and continually improve the efficacy of marketing. Effective response models will help identify which patrons are most likely to respond to an offer, and in turn to which offer patrons are most likely to respond. There are at least three main uses of response modeling that can improve marketing results:

1. Identify the likelihood of patrons to respond to the offer
2. Identify the offer(s) to which patrons are most likely to respond
3. Predict when a patron is likely to return

IDENTIFY LIKELIHOOD OF RESPONSE

The following example can help to illustrate the ways response modeling and optimization can improve marketing campaigns. Every year there are 1,000 discount tickets for a headline show that are purchased by the marketing department to distribute to players through a direct mail campaign. Last year, an offer for a complimentary room and 2 show tickets was sent to 100,000 patrons. All the tickets were accounted for within a week of when the offer was sent out and many guests were unable to use the offer because all the tickets were given away. This year the

Patron Analytics in the Casino and Gaming Industry: How the House Always Wins, continued

department wants to be more analytical about the tickets to make sure they go to the best patrons that will make the most money for the department. A response likelihood model is built to predict the patrons that are most likely to respond to the show ticket offer. This allows the department to optimize the campaign to drive the most profit from the best responders. Additionally, this can help to limit patron frustrations with offers becoming sold out or unavailable.

A response likelihood model can lead to lower mailing costs by identifying patrons that are very unlikely to respond to a particular offer. In the previous example, the business can identify the likelihood of response from all eligible patrons. After that, they can identify the most valuable patrons that are most likely to respond. This allows the business to estimate the expected response from the most valuable patrons and eliminate mailing(s) to the patrons that are of lower worth and/or are unlikely to respond. If the initial list of patrons does not account for all the tickets bought, then the lower worth/high likelihood to respond patrons can be contacted. Rather than sending out the blanket offer to 100,000 patrons as in the year before, only 50,000 patrons are mailed and the offer is redeemed by higher worth patrons. Occasionally, response likelihood models will lead to easy decisions, such as cutting out low worth patrons with a low likelihood to respond. However, more complex situations might arise since response models are never perfect. No matter how good a model and historical data are, there is always a chance that a patron identified as unlikely to respond will respond. Thus, when making a decision about patrons identified as unlikely to respond to an offer it is also important to balance that likelihood of response with the potential return on response. For instance, a segment of patrons is identified as being the least likely to respond to the offer in the example above. However, this is a small segment of patrons with a high predicted worth, and if just one of the patrons in this segment responds the entire cost of the mailing will be covered (in addition to other expenses). Instead of cutting them out completely, it is decided to mail the segment anyways because the upside is so high.

SEND THE OPTIMAL OFFER FOR EACH PATRON

We can generate likelihood scores for a wide array of offers to identify the best offer that the patron is most likely to respond to. For example, consider a patron who is predicted to have a high likelihood of response to the show ticket offer above. We also know the likelihood that a patron will respond to (1) a room only offer, and (2) a room offer with a complimentary dinner for two of equal cost to the show tickets. This information can be used to optimize the offer that is sent and the expected return on the offer. If the patron is predicted as equally likely to respond to all three offers, the offer with the lowest cost (i.e., the room only offer) should be sent. In a business where there are a variety of reinvestment outlets (i.e., rooms, gaming, retail, entertainment), response modeling can help identify which offer has the best chance of driving a response. If this is done for a segment as a whole, we can maximize response and profit by ensuring that the most expensive offer (in terms of reinvestment) goes to the patrons that might be less likely to respond. It might be tempting to offer the guest the option to choose which offer they want (especially if the offers are of equal value), but there is evidence to suggest that narrowing down a guest's choices is actually better and more profitable. In this case, it's a great idea to test whether a targeted offer based on a likelihood model performs better than an offer that gives the patron a choice between the available offers. Additionally, survey or preference data might be available that informs the business of what type of offers patrons are interested in receiving. Patron preferences can also be used in response models to determine whether a stated preference is actually a good indicator that a patron will respond to an offer targeted towards those interests.

A similar model might be built to identify patrons that are likely to only come over weekends and holidays (sensitive to convenience) or are likely to come whenever the best price is available (price sensitive). This type of information is helpful for building offers that drive patrons sensitive to specific periods. For instance, we know that patron A and patron B have similar levels of worth. However, patron A is equally likely to make a midweek trip or weekend trip and patron B is most likely to make a weekend trip. Given the limited availability of rooms, we might want to send patron A an offer with midweek availability only and patron B an offer with weekend or midweek availability. That way, patron A doesn't take a weekend date away from patron B. By doing this, we can maximize our response and ensure that we are filling the most number of rooms with the best guests possible. Furthermore, if patron A doesn't respond to the midweek offer we can send that patron a weekend offer the next time weekends are available (or if weekends are still available after the initial response spike from the weekend only offer).

Response optimization strategies can be extended beyond making financial decisions about what offer to send and whether to send an offer based on response likelihood and reinvestment. There might be no easily identifiable financial benefit associated with whether a guest is given the option to redeem a table offer or a slot offer with the same level of reinvestment. Although common sense tells us not to send an offer for table games if the guest has exclusively played slots for 10 years, it might be difficult to determine whether the choice of offer has no effect, improves response, or worsens response. By identifying the offer(s) that guests are most likely to redeem it shows that the business is paying attention to what a guest likes and personalizing offers for them based on their behaviors or interests. Below are other examples of common scenarios in the gaming industry where marketing optimization can be used to personalize offers and improve response:

- To identify a segment of patrons that like to play a certain brand or style of slot machine for a marketing campaign for new machines featuring that brand or style is coming to the floor

Patron Analytics in the Casino and Gaming Industry: How the House Always Wins, continued

- To personalize marketing with offers, pictures, and content based on the restaurants, shops, games, and shows that patrons have attended or frequented
- To theme a slot tournament around a theme or brand that is popular among high end players

IDENTIFY WHEN A PATRON IS LIKELY TO RETURN

In addition to having some information that helps determine to which offers a patron is most likely to respond, it would be nice to know exactly when a patron was planning on making their next trip. Although we might not know exactly when a patron is likely to return, fortunately we can make a pretty good prediction about it. There are a variety of methods that range in complexity that can be used to assess when a patron will return, including **frequency analysis, regression, and survival analysis**. Knowing when a patron is likely to return is beneficial as it helps to identify patrons that haven't made a trip in the expected amount of time and are at risk of leaving. First, the business needs to have an idea of the average or median time between trips. This might need to be segmented based on geography, worth, or even historical frequency. Patrons that have not made a trip within the decided amount of time for their segment are subsequently flagged and dealt with appropriately.

Historical data can help to identify segments of patrons that are expected to make trips weekly, monthly, quarterly, annually, bi-annually, and so forth. Marketing can integrate information from predicted worth, optimal offers, and time to next trip to maximize campaign success in a number of ways. The business can save money by adjusting the frequency of offers for patrons that are not identified as likely to come back for longer periods of time. Instead of sending the patron monthly offers, they can sent quarterly offers with longer valid windows that allow more time to book. Or, for example, if the patron only comes annually around his/her birthday, we might only send an offer annually around the patron's birthday. Conversely, campaigns might be created with the goal of increasing the frequency of visits from higher worth patrons. Casino marketing should have the goal of generating trips sooner than expected and converting patrons into more frequent visitors. Additionally, time to next trip analysis can be used to identify when it has been too long and the business is at risk of losing the patron. In this case it might be useful to send an offer using "last chance" "we miss you" messaging. The offer might also need to be slightly better than what the guest has received in the past. By knowing when a patron is likely to return, we can adjust marketing strategies appropriately in order to save money on mail costs, retain guests, and increase loyalty.

SAS SOFTWARE AND PATRON ANALYTICS

To fully utilize the power of predictive analytics it is a good idea to make sure and use software applications that can handle large amounts of data. Although basic statistical packages can be used for many of the statistical analyses, more powerful data manipulation software, such as SAS Enterprise Guide® software with SAS/STAT®, is recommended for large sets of data used in casino analytics. SAS Enterprise Guide enables you to do almost all of the data manipulation, exploration, and mining techniques that are mentioned in this paper. Additionally, various statistical techniques are available through SAS procedure, such as the following:

- The PROC REG and PROC LOGISTIC procedures are used for linear and logistic regressions respectively
- The PROC DISCRIM procedure is used for discriminant analysis
- The PROC LIFETEST and PROC LIFEREG procedures are used for survival analysis

Some functions, such as decision trees, are only available in the SAS Enterprise Miner™ software. SAS Enterprise Miner is ideal for building many models, as it significantly decreases the amount of time and energy required in the data mining process. Additionally, SAS Enterprise Miner allows the user to quickly compare models against one another using visual charts (i.e., cumulative lift) and fit statistics to determine which model fits the data best. For instance, the model comparison function can be used to compare a Decision Tree to a Logistic Regression to determine the model that is best for predicting responders. SAS Enterprise Miner is an excellent tool for the analyst looking to quickly and easily build multiple models and implement them into marketing processes.

CONCLUSION

A solid casino marketing foundation built around patron analytics will lead to improved marketing results by identifying the best patrons and the best ways to drive business from those patrons. This paper has supplied a number of approaches to the most common questions in casino marketing. All it takes is some creative thinking and hard work to uncover actionable information from the vast amounts casino data. This information can be used to improve response, revenue, and profit driven by marketing campaigns. Furthermore, advanced analytics can be used to generate powerful models to be used for the prediction of future behavior.

The good news is that the field of patron analytics in the gaming industry is still growing. Many casino loyalty programs are focusing on using analytics to drive more than just gaming revenue. This shift will lead to additional data about non-gaming behaviors and the integration of data from non-gaming sources with gaming data. Additional

Patron Analytics in the Casino and Gaming Industry: How the House Always Wins, continued

sources of information and revenue are starting to be integrated, such as feedback and data from social media. The potential legalization of online gaming could provide exponentially new amounts of data and create wide open territories for patron analytics and marketing integration. Ultimately, the goal of casino patron analytics is to provide actionable results faster that can drive offers in real time and build stronger relationships between the casino and the patron.

REFERENCES

Las Vegas Convention and Visitor Authority 2010. "2010 Las Vegas Year-to-Date Executive Summary"
<http://www.lvcva.com/getfile/624/ES-YTD2010.pdf>

RECOMMENDED READING

- "Discriminant Analysis, A Powerful Classification Technique in Data Mining"
<http://www2.sas.com/proceedings/sugi27/p247-27.pdf>
- "Survival Analysis using SAS®: A Practical Guide" by Paul Allison
- "Combining Decision Trees with Regression in Predictive Modeling with SAS Enterprise Miner"
<http://www2.sas.com/proceedings/sugi30/074-30.pdf>
- Berry, M.J. & Linoff, G.S. (2004). Data Mining Techniques. Indianapolis, IN: Wiley.
- Kilby, J., Fox, J., & Lucas, A.F. (2004). Casino Operations Management. Canada: Wiley.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Scott Sutton
The Venetian and Palazzo Resort Hotel and Casino
3355 Las Vegas Blvd South
Las Vegas, NV 89052
(702) 414-4748
scott.sutton@venetian.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.