Q&A

NICK ROSS

DIRECTOR OF ANALYTICS

SEGA®

Game Analytics & Business Intelligence

Nick Ross, PhD, is Dr. Data at Sega, where he leads a team working on business intelligence and user acquisition. Before Sega, he was Director of User Acquisition and Analytics at TinyCo and has spent time employed at Bates White, an economic consulting firm, where he answered math, statistics and economics questions for lawyers. Nick holds a PhD from UCLA's Anderson School of Business, where he researched empirical and theoretical issues in economics, finance, and accounting...

Nick, from your perspective at Sega, where have you seen a real impact on ROI when it comes to investing and applying data analytics?

The most obvious area that our data analytics system has generated ROI is in user acquisition – being able to look end-to-end to effectively understand where our marketing dollars are generating value and where they are not. The way most user acquisition works today is on a cost-per-action model, such as paying per install or click. The simple question is, on aggregate, if I purchase 1,000 users at a buck a piece, is my expected return on that greater or less than a thousand dollars? Is that going to help me hit my business goals? When I speak with my boss, that's the first question he asks because it's money directly out the door.

There's a cottage industry popping-up to assist developers' monitor the effectiveness of their spending. Companies like Ad-X, HasOffers and Kochava are integrating with adverting partners to track ROI for developers. From my perspective, there is value in building some of this expertise in-house, as we are doing at Sega, but if you're a smaller developer you can outsource all this.

When it comes to actually choosing vendors, is there a way to determine what works and what doesn't? How do you go through the process of separating the good from the bad?

There's certainly value in determining which data tracking vendors are worth investing your time and money.

Big data collection is very quickly becoming a commodity. Even just a few years ago, if you wanted to collect a mass amount of data about your users' buying habits you were left with a bunch of tools that didn't provide granularity. The data provided was aggregate, such as the number of people completing the tutorial or interacting with a specifically tagged event.

These services are still available, but what we've seen is that the costs, in time and money, of this type of work has dropped. It has become pretty easy to collect this information in-house and gain access to the underlying, much more detailed, information that the last generation of tools failed to provide.

Previously, you would expect to spend tens of thousands of dollars a month on analytics, mainly on the collection system. Now you can bootstrap that for a couple of grand a month – and if you're really cutting back, you can probably collect your own data for a few hundred a month, even with a high DAU game.

So companies providing those services are feeling price pressure and are working on providing a bit more insight into the underlying analytics, beyond what traditional collection tools provided. But, at the end of the day, who knows more about game than the game designer? I'm not sold on outside services abilities to *interpret* the data from a game that they didn't design and are not, like my game designers, thinking about 24/7.

Obviously the ultimate point of data analysis is to de-risk – to know the finer detail and then act upon those findings. But are there inherent risks in actually undertaking the process of data analysis?

Yes, there are. Collecting data is one thing and analyzing it is another.

Obviously, the biggest risk is that the decision you're making is based on data but which is still an uninformed decision. By 'uninformed' I mean that you may not fully understand the information that was presented. For example, there may be details about the data collection process that you are not familiar with, which effects its interpretation of your reports. Consider that some data collection services don't buffer and queue events on the client, so they only state events that happen when the game is connected to the internet. This creates a bias. Whether this bias means much is dependent on your game and the decision you are making. Ultimately, being uniformed in this manner can lead to a bad decision.

➡

Over the last decade there has been a true renaissance in the collection of in-game user behaviour data. While there is a clear recognition of impact of analytics on the games-as-a-service model, the use of this data in pre-production design is still in its infancy. At this year's Game Analytics & Business Intelligence conference, Nick Ross will cover several case studies on how design and analytics can work together in a game's life-cycle to maximize the potential for success. Before the event, he spoke with *Gaming IQ* about his views and experience…

'We're getting rid of the SQL monkeys… There are now solutions for many of the technical pieces across a wide spectrum'

Just how common is it that studios don't understand their own analytics? Is it an area that's seeing progress when it comes to remedying those pitfalls or does this remain a widespread ailment within the industry?

In my mind, the biggest failure for studios not understanding their own analytics is a failure to understand how the information gets to them. This results in a credibility failure; which is what I see more frequently than a lack of understanding.

We tend to say that the (data) pipes are leaky. The question is: how leaky are they? There's always going to be some errors that you're not catching, some exceptions, some events you're missing. That's always going to be the case. But it helps to find out when and why you're missing those. When I was at TinyCo I always made the joke that our data was a 99.9% biased sample,

'It's important to keep in mind that cross-platform biases exist and how data generating processes effect the metrics'

and that was absolutely the case. I knew we were losing a tiny percent of our data and I knew exactly who we were losing, so we were clearly biased! But at the same time it was still a 99.9% ample. Merely understanding that allowed us to make better decisions and, more importantly, understand the limits of what decisions our systems could make.

When you're using third-party providers, one of the issues I see is that people may not know what a number is. How much do you know about the sample that created that number? Do you have the full picture as to what's breaking when? And how is that affecting the information you are seeing? When you make decisions based on this information, are you taking all that into proper consideration?

From my perspective, that's the most common error. The other failure that I commonly encounter is smart people using information vindictively. That's a very, very common problem. Rather than looking at data to improve a product, they're looking at data to make a point, or they're looking at data because they want something. These types of behaviours are very common and have a very negative impact on an organization and the use of analytics within that organization.

That's something that you have to believe is culturally unacceptable and then knock it on the head when you see it happen. To help this, one of the things I like to do is to gamble on our data. Instead of me forcing the issue with the team, I'll say, 'All right, you know what? I think I'm right about this particular thing. Let's put $10 on it!' You know, not a lot of money, but enough that everyone takes it seriously. I find that that's a really good method, personally, of making sure that I don't fall into these behaviours – and nor do others.

Changing the culture of your colleagues is an interesting point. What about further up the chain when it comes to getting buy-in from those who are in charge of where the investment is going? Is it ever much of a problem to push money towards better

analytical processes and to justify those expenses? And is there anything you would recommend for start-ups that may be faced with tougher investment decisions like these?

I think for a lot of game developers, it's not going to be the issue of money most of the time, it's going to be the spot in the sprint to handle the tracking issues, especially with everything now being so VC-fuelled. It's getting enough time to be able to take care of the issues you find. So, at least from my perspective, it's often less about the money.

When I got hired at Sega, the analytics budget

➡

OFF THE RECORD

Any hobbies?

Surfing, sailing and mountain biking.

Tell us an interesting fact about you

My PhD is in Accounting.

Top 3 favourite games?

1. Secret of Mana
The play, music and story were perfect.

2. Legend of Zelda
A link to the past...

3. Borderlands
I loved the humor, graphics and RPG-ish elements.

was, if anything, a bit too high. The amount of money needed to run what I consider to be a successful analytics pipeline, to employ analytics tools, to set up successful tracking, it all costs a lot less than what people expect. Part of it is that I'm willing to build some of this in-house and I already know how to avoid some common pitfalls. But part of it is just people expect. Part of it is that I'm avoid some common pitfalls. But part of it is just that it used to cost around $1,000 to get a DB server on Amazon and now it's a couple of hundred bucks a month. With computing prices falling as rapidly as they have been it is easy to end up with an overfunded analytics budget.

So it's generally not about the money, but about development time. It's about having a new game, finding that one of the events doesn't fire properly, needing a dev to spend a day in the code to figure out why not, and convincing PMs to spend time on it. That's much more expensive, especially when your PMs have been designing sprints and they're thinking 'Okay, I have a two-week sprint (or whatever type of coding review process). I have 20 man-days and I'm going to do two man-days on this thing, two man-days on that thing, one man-day on this thing, and now you want me to take time away from fixing some UI/UX to make sure a tracking event is firing correctly?' I think that's a much harder sell..

One of the hot-button issues we're looking at this year is data visualisation and how people in your position can specifically take a block of information and get it to make sense to shareholders who may be less technically inclined, or to various parties with differing concerns. What seems to work for you? And what's the outlook for options in the future?

The big trend that we've seen – and I hope will continue to see – is that we're getting rid of the SQL monkeys. When I started in "big" data, you had to have at least one or two people whose job it was to maintain the webpages, make sure that the SQL was doing whatever it was supposed to be doing, making sure the JavaScript visualisation layer was working and that type of thing. We've
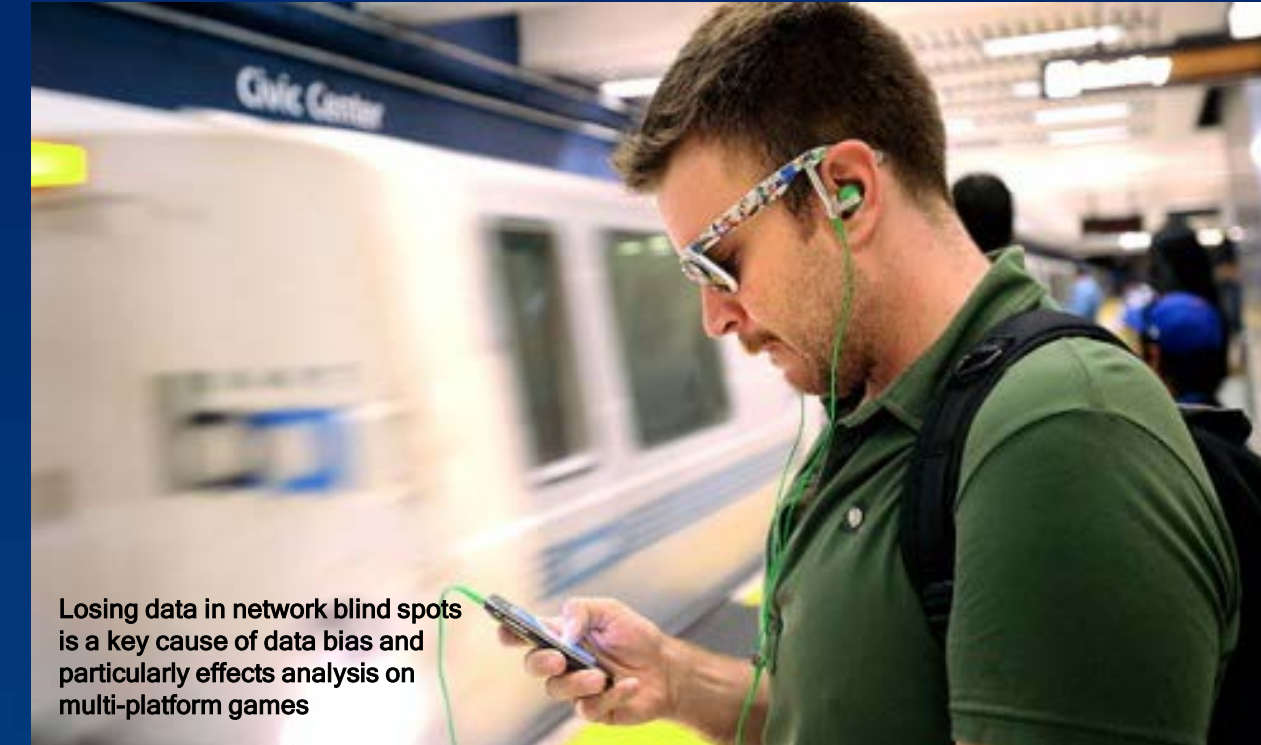
Losing data in network blind spots is a key cause of data bias and particularly effects analysis on multi-platform games

come a long way from that. There are now solutions for a lot of the technical pieces across a wide spectrum of use cases.

Do you want to have a bunch of SQL monkeys or code monkeys working in JavaScript and D3 to make really pretty visualisation? That's really expensive in both time and money. Nowadays, there are things like Looker and Tableau, which – if you have a basic understanding of SQL – can give you 90% of the answers you're looking for. Those tools on the packaged side of that spectrum have come a long way and are becoming more and more commoditised, which means that the prices on them are falling very gradually.

Visualisation is much more prevalent today, it's much cheaper, and the types of people you need in order to have good visualisations are a lot less technical than they were five or ten years ago. When I hire data analysts to make visualisations, I basically look for smart recent grads. I don't look for a coding background. If they've got a good head on their shoulders and can understand basic math or maybe took a couple of calculus courses – that's all I need. A few years ago, if I wanted to do the same thing, I needed much more than that. I needed someone who at least could

understand a bit of programming. That's no longer the case.

So where do you see that evolving in five years from now, if at all?

I'm hoping that the stuff on the non-packaged side of the visualization spectrum, the very technical stuff, just disappears. I expect the tools to get simpler and simpler to implement as time goes on so that my team wastes less time futzing with software and more time actually looking at data and making decisions on their basis. I absolutely see that trend continuing.

Take Tableau, for example – which I think is an excellent tool. While it has some extremely annoying shortcomings, but other next generation tools that are out there, a lot of which are in the field of data, are very powerful. I like them a lot. And they're web based. You can use them on your mobile phone. You actually don't need to spend $1,000 a user, like you do with Tableau, and I'm very excited about that.

➡

What about cross-platform disparity? How tough are you finding it to work across different environments – mobile, social, video, and so on – and maintain a consistency in the information coming through the pipe?

Gaming platforms all come with their own little tricks and their own little differences. Let's take the following experiment: You can play Candy Crush Saga without being online. If you turn off your internet or put your phone on airplane mode, you can still play Candy Crush. But when you're on the Facebook platform, you pretty much always have to be online to play.

When King.com is collecting information, they basically have 100% of the events that people play on Facebook. They have excellent tracking and they're losing very few events. Their KPIs can be extremely nuanced and they don't need to worry about much bias. Yet on the mobile platform, the only events they're going to be able to collect are when the user is online. So either they're buffering – putting the offline events in a queue to send later – or they're just tracking things when the users are online. Either way, biases and differences will occur because of that and one platform's data won't look perfectly like the other.

I live in Oakland and I commute to San Francisco. Every day I take the BART [Bay Area Rapid Transit] and it goes underwater. And every day I see people, as soon as we enter the tunnel, when they lose signal, pull out their phone and starts playing games. When we pull into the other end of the station, they put their phone down and they go back to email. You're not tracking that user.

So naturally, we see these cross-platform nuances occurring. Continue that thoughtexperiment and now through Activision into the mix, which has Call of Duty, with its popular online multi player component.Because the game is on console and requires an online connection Activision can do amazing tracking in this part of their game, but in terms of being able to track the single-player experience, the only people that they're tracking are the people who are connected to the internet. Now those people are probably the very good players, the regular gamers. If you're a complete noob and you've never played a first-person shooter before, chances are your PS3 isn't connected to the internet.

In other words, in this example you're often not seeing the information of the less sophisticated users. You're never going to see where they're dying. They could just be getting trapped in the same spot over and over and over again – as noobs do – but you're going to have a hard time finding that spot because the data is being lost.

Just in that simple thought experiment, comparing a game on Facebook, mobile and console you can see that there's a lot of potential difference in the reported metrics. In terms of maintaining a consistency across platform, I think that the important idea to keep in mind is that theses biases exist and to make sure that when you make decisions, you understand how the data generating process effects the metrics you're looking at.

We talk about using analytics effectively and most of the discussion out there seems to be around how to apply that analysis in order to adapt to the market. But should there be

Sega's most anticipated new title, *Alien: Isolation*, hits shelves in October

'I'm going to be looking to find out how other data analysts tackle these problems, how they organise themselves, and whether there are best practices I can take from that'

more focus on taking those numbers and looking to actively change customer behaviour?

In my mind, you have to do both.

Let's take Candy Crush as an example, again. Say you have the game designer. He designs this game and he expects certain behaviours in the game. He makes the first couple of levels very easy to win because he wants the success rate on those levels to be very high – say 95%. Then he wants level nine to present a little bit of a challenge, so he anticipates the success rate to be on the order of 60%. Then, perhaps right after that challenging level, he gives the user another couple of easy levels. You can make a curve like that over all the levels. Now you've put the game out there and you're collecting data and you can calculate those success rates. In my mind, the first thing that you need to do is tune the game to what the game designer has laid out because until you do that, you don't actually know whether the rest of the designer's predictions are accurate.

Once that occurs, only then can you start talking about shaping behaviour, and more nuanced monetisation, using tools such as A/B tests. What frustrates me is jumping to the second stage without taking the time to tune the game to the designer's specifications. Look, I obviously don't expect a designer to have every number calculated out to 3 digits, but a good designer will have thoughts about the three most important numbers, or what the trends would be on those.

If it's Candy Crush, I'd think the designer would say things like 'I want there to be gate levels (high difficulty) and I want the gate levels to be at level 12 and level 45, and I want the average person to have to try 4 times before they get through." That's the level of detail they would need to create an initial tuning. I'm not expecting a designer to distinguish between 7% and 8%.

When you think of a player being frustrated, what does that frustration look like? When you think of a player having fun — what does that look like? A good game designer will have thought about that and a good development team will line the game up to the designer's original vision.

The question then lies in how we want to change that if the game's not monetising or, simply, not fun to play. But you have to get it to the designer's vision before you start retooling.

Looking across the vast timeline of gaming history, there seems to be a shift when it comes to gamer expectations on the product. Years ago, it would be common to play a game and never complete it. Nowadays, I don't imagine most people would buy a game without expecting to complete a core mission or goal. Some of those players set out to do it within 24 hours. How much does this represent a shift in the habits of the end-user and is it impacting buying decisions? Do you spend much time cross-referencing data with behavioural analysis or wider market and cultural trends?

I think it depends on the player and the objectives. Bigger companies that have a higher budget for this type of thing do a lot of user research. A lot of game design that we do here is based in what we call 'profiling', where you or the game designer sits down and assesses the types of players who are going to play this game: Type A, Type B, Type C, and Type D. These are the other games that Type A enjoys; these are the ones Type B would enjoy; this is Type A's playing pattern; we expect the market of Type C people to be big enough to actually support the budget of this game; and so on.

If you have a bigger budget, once you've done all that, you can of course then go and find those people and make them play the game before the launch. You can undertake marketing studies with tests and surveys. You can cross-validate whether the market is going to support that type of title. You would test a game like Skyrim by finding people who would normally play a Skyrim-type game, put them in a room with the product and find out if they actually like it. But all that can get really expensive really quickly. Bigger studios have the luxury of doing that – for others, it's a tough ask.

You'll soon be discussing many of these issues at the Game Analytics and Business Intelligence conference here in London. Aside to bringing a lot of value to our speaker panel, what are you yourself hoping to draw from the event? Is there anything to which you're looking for an informed answer?

For me, the value in these types of events is in understanding how other people are thinking about these problems. As analytics people, we're not often going out to conferences and we're not having lunch with each other that often. So just getting the thoughts of others doing this same work is really valuable.

A lot of other conferences aren't even geared towards the analytics side. They tack a little bit on, but because others in the room aren't focused on this aspect of game design, the dialogue can't be technical. They can't talk about the forefront of analytics. They can't really call into question the expectations for the years ahead or how we expect to use numbers in the foreseeable future. A talk like that would never be approved by

GDC because it's too technical and doesn't appeal to a broad enough audience. The Game Analytics and Business Intelligence conference can do that.

My background was in a small gaming start-up and now I'm at Sega. I don't have the experience of others who have perhaps come from larger games studios, or who have invested more into analytics. I don't have 35 data analysts. So I'm going to be looking to find out how they tackle these problems, how they organise themselves, and whether there are best practices I can take from that.