

# We're Not Listening: An Open Letter to Academic Game Researchers

By John Hopson

[\[Author's Bio\]](#)

Gamasutra

November 10, 2006

[http://gamasutra.com/features/20061110/hopson\\_01.shtml](http://gamasutra.com/features/20061110/hopson_01.shtml)

## Introduction

Academic interest in games has risen quickly over the past decade, but the games industry has never shown a similar interest in academic work. Every year there are books, journals, and conferences dedicated to studying games and how people play them, but most games professionals never read this work nor attend these conferences.

Sure, individual designers, producers, and developers listen, but the industry as a whole has ignored an entire field of study dedicated to studying it. There have been academic panels at industry conferences, but the vast majority of the conference attendees have walked right on past. I've sat in these sessions and heard the researchers vent their frustrations: They're doing wonderful stuff, why won't the industry just listen?

We have scores of smart, professional academics out there doing great work, learning and thinking new and fascinating things about games every day. This work should be having a huge impact on the games industry, the kind of impact that university researchers in chemistry or computer science have on their own fields. Why aren't we seeing those sorts of breakthroughs improving our ability to make better games?

I often refer to myself only half-jokingly as a "recovering academic." I received my Ph.D., did my research, wrote journal articles and grant applications, taught undergraduates, and attended academic conferences so dry that any normal human being would have died of dehydration. However, on the brink of diving into a full-blown academic career, I wrote an [article](#) which drew some attention from the industry and I jumped at the chance to impact games directly.

For the past several years I've worked as a professional games researcher on major titles such as *Halo 2* and *Age of Empires III*. I've been in the trenches, doing applied research on how people play games and then have been working directly with development teams helping them to use the results of that research to make real changes happen in real games.

But honestly, even I walk past most of the academic presentations at industry events. Even I have trouble really getting excited about most of the games research being done out there. From the perspective of someone on the inside, the average piece of academic games research just doesn't get the job done. It's not a question of the quality of the research or the intelligence of the researcher or the game makers; it's a question of bridging the gap between the academic and business cultures.

For the purposes of this article, I'm going to make a couple of assumptions. First, this article is for researchers who want their ideas to be taken to heart and implemented by non-academics in the games industry. Pure research and theory are beautiful things, but they're outside the scope of this article. I'm not saying that academics have to care about what the industry thinks of them, but for those who do this is the best advice I can give on how to make sure the industry takes your work seriously. Secondly, I'm going to assume that you've already done your research and have the findings ready to go. This article is about the final mile, going from finished research to real implementation in a shipped game.

## **Rule #1: Return On Investment**

People in the games industry make games because they love games, but it's still a business. Every person working on a game can justify their existence on that project in terms of their impact on the final product. This is not necessarily in strictly monetary terms ("If we hire another tester, we'll up sales by 3%.") but generally more in terms of final game quality. ("Our testers are stretched too thin, if we hire more we'll catch more bugs and ship a better game.") Every employee, every computer, every meeting, every stapler in a game studio is a tool for making the final game better. If games researchers, inside or outside the industry, want to be taken seriously, they have to justify themselves and their work the same way.

When a researcher presents a product team with a set of research findings and recommendations, they are asking the team to invest time and money implementing their proposal. In order to convince the audience to spend that time and money, the researcher has to show clearly how that investment is going to pay off. This needs to be something beyond "this will help players identify more strongly with the main character".

The researcher must lay out the entire impact of the idea, from the cost of implementing the proposal to the resulting changes in player experience and the metrics for measuring that impact. Getting players to identify with the main character is great, but researchers have to finish the rest of the sentence: "This will help players identify more strongly with the main character *which will result in an improvement in measures of overall player satisfaction and an increase in total playing time.*"

By the way, if the research doesn't include specific practical recommendations or a measurable impact on the final product, don't bother trying to sell it to the industry. From the average industry professional's perspective, there are only two things of value being said in a research presentation: the recommendations and their predicted effects. Everything else, the background research, the brilliant theoretical breakthrough, the clever development of the ideas, falls on industry ears like the "wah wah" noises made by Charlie Brown's teacher. I'm not saying that qualitative or theoretical work isn't worthwhile; I'm just saying that the industry is generally not going to listen.

## **Rule #2: Speak The Language**

Academic writing is abnormal. I know that by the time you escape grad school the rolling cadences and ritualized forms of the journal article are graven on your very soul. But really, you might as well present your research in the form of an interpretive dance as hand a producer an article written for academic publication. Reading an academic article is an obscure and highly specialized job skill, one which most of your potential audience doesn't have the time or desire to learn. It's up to the researchers to make their work accessible to the audiences they want.

The goal here is to convey information clearly and easily to specific types of game industry readers. They're already going out of their way to read it, so make it as easy as possible for them to assimilate your proposal.

This is not to say that researchers need to dumb anything down! This audience is just as smart and critical as any journal review panel, and it's the one who is potentially going to have to spend days or weeks implementing the ideas. Rejecting the proposal is the easy route, the one that doesn't mean changing the schedule or undoing any finished work.

Think of it like grant writing. The way research is framed in a grant proposal differs from the way it is presented in a classroom because both the audience and the goal of the presentation are different. The core of the research remains the same no matter what the audience, but different aspects are important in different forums. The same talk that knocked the audience's socks off at CHI is going to fall flat at GDC. Don't write a research paper, write a business proposal.

A couple specific guidelines:

- **Start fresh.** You're talking to a new audience, so you should be creating a new presentation. Don't just take out the same PowerPoint presentation from the last academic conference and adapt it. Start a new slide deck and add every slide with a view towards the new audience and what it needs to know.
- **Lean and Mean.** Think one page. Academia rewards writers for covering all the bases, making sure every tiny step is supported and documented, and for detailing how the current work fits into the existing literature. This is good and appropriate for academia, but industry readers don't care about most of that. Start from the bare bones of the argument and put on the minimum flesh necessary to convince your audience to invest in your ideas.
- **Use examples from bestsellers.** A good example from a popular game is more effective than a great example from something they've never heard of. Industry people often suffer from an "if-they're-so-smart-, why-ain't-they-rich" attitude towards smaller titles. Even if the small title is a perfect example of how the theory works, they're going to be less likely to listen if they haven't heard of the game ahead of time. Commercial success is one way of making sure that the audience will respect your examples, but you can also use titles that are well known or critically acclaimed but which weren't necessarily huge blockbusters. It's also important to keep your examples as current as possible, because many industry folks will see a three-year-old example as ancient history.
- **Look forward, not backwards.** Lose the lit review. Don't quote references. Don't worry about background material. This is about specific, concrete recommendations and the impact on their game. Shape the presentation from the recommendations outwards, providing only the background absolutely necessary to justify the recommendations.

### **Rule #3: Smaller, Faster, Cheaper**

Working in the games industry can be brutal, involving fast-paced schedules and eighty-hour work-weeks at times. The people listening to your talk already have a full workload. They've already been cutting features to make their production milestones, often features that

represent some of their best ideas and strongest held beliefs about games. And then there's this academic who's never shipped a game standing up there telling them to rip out weeks of work in order to implement some pet theory. Give us a break!

Furthermore, games work on a multi-year development cycle. For a given article or conference presentation, the audience members will be spread across all possible phases of development. Some will be working on games still in the prototype phase, some will be in full production, and some will be in the final crunch to release. If the recommendations involve major changes to the design, the odds are that the majority of the audience has already passed the point where those sorts of changes can be made on its current project.

All recommendations are not created equal. Some are going to be easier to implement than others. Making a change, any change, to a game that's already in production is akin to doing auto repairs on a car going 60 mph down the highway. The very fact that your audience is even listening to your recommendations, even *considering* making any kind of change, is a major concession. It's up to the researcher to choose recommendations that make it easy for the developer to say "yes."

The best recommendations should be...

- **Feasible for one person to implement.** It's much easier to convince one person than ten. If a suggestion can be taken to heart by one designer who can then turn around and implement it in just his own area of responsibility, that's a lot more likely to happen than a change that requires the entire design team to agree.
- **Scalable.** A good change is one that can be quickly implemented on a small scale for trial purposes, and then expanded if testing shows that it's working. Make the required leap of faith as small as possible.
- **Modular.** The fewer aspects of the game that need to be changed, the better. The more a suggestion cuts across boundaries (art, design, UI, etc.) the more likely it is to be ignored. Changing either the appearance or behavior of a character is easier than changing both together, for example.
- **Parametric.** The easiest things to change in a game are things that can be represented by single numbers. Changing what a gun looks like is hard. Changing the damage it does by 10% is easier. Whenever possible, pick recommendations that can be implemented by changing individual game parameters.

#### **Rule #4: Prescriptive, Not Descriptive**

One major difference between academia and industry is that academic work can be purely descriptive and still be successful. Discovering and describing a new phenomenon can be a genuine academic victory in itself. However, while understanding games is a research discipline, actually making games is an engineering discipline. There needs to be a clear and explicit path from the imagined beauty of research ideas to the ugly reality of implementation.

One particular area where understanding falls short of implementation is in the realm of nomenclatures and classification systems. In academia, developing a way to classify players into specific defined types can be a very real and present help to future research on the topic.

However, to make that research useful to developers, it's important to take the next step and give concrete examples of how classifying one's players helps to make a better game. Ok, so

we now know that 10% of our players are "Type 3a explorer-monkeys." Now what? Does that mean 10% of our content should be exploration-focused?

Descriptive statistics is another research area where it's important to take the leap to prescriptive recommendations. Knowing that most players on an MMO have at least five characters is nice, but not everyone is going to make the leap to the idea that players should have a universal alias that lets them talk to their friends regardless of what characters they're playing at the moment. (A very nice feature found in *City of Heroes/Villains*) The specific recommendation that grows out of the statistics needs to be the point of the article/presentation, not left to the reader or tacked on at the end.

## **Rule #5: Prove It**

When I first told people I was going into the games industry, every single person I told had an idea for a game I should make. Every player of every game has recommendations for what needs to be changed to make the game better. Developers are buried in suggestions, anecdotes, and feedback, and almost all of it is shortsighted and badly thought-out.

The way to ensure your ideas stand out from the cacophony is to support them with a level of proof and certainty that the average wannabe off the street can't match. The best way to do this is to provide an active, working example of your recommendations, with numbers to back it up. Pick a currently popular game and put your idea to work. If your theory makes predictions about FPS level design, grab an existing game level off of the Internet and modify it to fit your recommendations. A side by side comparison with statistics to demonstrate the impact of your changes is the most powerful tool you can have to convince an industry audience.

Now, this example of modifying an FPS level requires a certain degree of technical skill, and that's important. People with technical skills get more respect in the games industry. Beyond that, it puts your recommendations above the level of the average raving fanboy who buttonholes a game designer and rants about why his favorite weapon/class/vehicle/strategy needs to be buffed. It shows them that this idea matters to you, and proves that you understand what it will take to implement the idea in a real game. One working model of your idea is worth a thousand eloquent, logical research articles.

Obviously, there are some topics that are very large and difficult to model, but almost all of those ideas are also too large for developers to really change mid-project. If your idea is too complex for you to test on any scale, it's probably too complex (and therefore risky) for anyone to actually implement in his game.

## **Rule #6: The Customer is Always Right**

A lot of research being done on games simply isn't relevant to the day to day work of the industry. It's not bad research; it's just focused on things that the industry doesn't particularly care about. Fortunately, there's a really straightforward technique for ensuring that your work is relevant to industry game developers: **Ask them.**

Ask them before you do the work. If you're doing a giant longitudinal survey of players in a particular MMO, contact someone at the company ahead of time and talk to him or her about your study. Start by contacting the game's community rep and explaining your project, they should be able to forward you to the right person within the company. They may be able to

provide you with internal data or to suggest lines of inquiry which might not have occurred to you.

Game designers don't release their games and move on to the next project; they watch their players and player communities with all the apprehension and passion of parents at a grade-school play. Having good research done about their game is beneficial for both the studio and the game, and most developers are happy to work with researchers when they can.

Be warned, however, that there are some potential costs to industry partnerships. You will almost certainly have to sign an NDA which will give them some say over what you can release and when you can release it. Furthermore, just as most researchers don't understand industry concerns, most of the people you'll be talking to in the industry won't intuitively understand what's involved in doing good research.

For example, they might try to draw conclusions from your data that you feel aren't scientifically valid. There are also potential ethical concerns unique to research collaborations with industry. If your MMO survey includes questions about practices banned by the game's terms of service such as item duping or gold buying, the question of whether or not to release those results to your industry partner becomes more than merely academic.

## Fair's Fair

Now that I've gotten all that off my chest, let me back off a little on two points. First, what I've said here isn't necessarily news to everyone. There are academics who have forged close ties to industry, working directly with developers to produce work that is equally successful as both pure and applied research. Others have found ways to present their information in snappy, bite-sized chunks for easy assimilation. One example of an approach that's on the right track is the "Game Studies Download" presentation from GDC 2006 (<http://www.avantgame.com/top10.htm>), in which scholars used short presentations of each piece of research and presented very clear implications from their studies.

Secondly, now that I've spent the majority of my time ranting at my academic brethren, I'd like to take a brief moment to turn my criticism on my industry coworkers. **Start listening.** Sure, academics sometimes appear to be ivory tower dreamers, but that doesn't mean they're necessarily wrong. In the day to day bustle and pressure of the industry, we simply can't try all the interesting ideas out there. The less focused atmosphere of academia can lead to some genuinely new ideas, things we might miss because we're too far down in the trenches.

Furthermore, **academic games research is free.** It's a rare game developer that can afford to pay for blue sky research, but these folks are happily doing this work for us. Even better than that, **academic research is public.** One of the few downsides of my job is that some of my work is (appropriately) considered to be trade secrets by Microsoft and the games studios. There are a number of companies out there doing really interesting applied research, but very little of that research will ever be available to benefit your game.

Remember Joy's Law: "No matter who you are, most of the smartest people work for someone else." Tapping into the work being done by academic researchers is a cheap way to put more smart people to work on your game.