The Four Myths of Game Design

ERNEST ADAMS | *Ernest is an American freelance game designer currently living in England. He was most recently employed as a lead designer at Bullfrog Productions, and for several years before that he was the audio/video producer on the MADDEN NFL FOOTBALL product line. He has developed online, computer, and console games for everything from the IBM 360 mainframe to the Playstation 2. His e-mail address is ewa@earthling.net, and his professional web site is at http://members.aol.com/ewadams.*

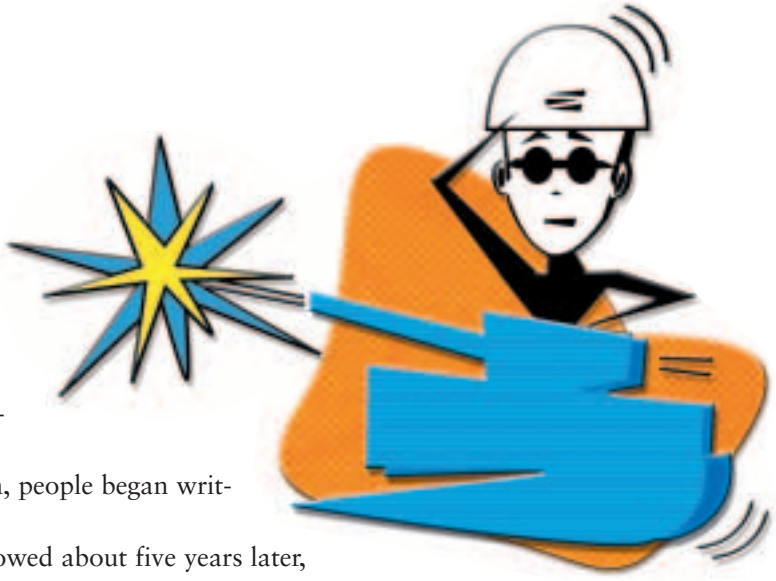

Thirty-five years ago, when time-sharing operating systems appeared for mainframes and the printing terminal became common, people began writing computer games. The game industry followed about five years later, if we count the earliest arcade machines. In the course of the last three decades, we've wandered down some strange paths, hit a few dead ends, and witnessed the evolution of a new entertainment medium, perhaps even a new art form.

We've also slowly evolved a discipline of sorts, a way of thinking about games and the people who play them. Much of this accepted wisdom is accurate, tuned by many years of trial and error and filtered by the natural selection of the marketplace. For example, we now know that great graphics alone do not ensure a game's success — the few games which violated that principle did not survive or reproduce. But a few bad ideas have managed to hang on as well, perhaps because they're not quite lethal enough to kill a company that relies on them. In this article, I'm going to discuss four commonly held beliefs about games and game design that are erroneous. If we could get these ideas out of the gene pool, we'd all be better off.

Myth #1: We Are Our Own Audience

The idea that we accurately represent our audience is the foundation of just about every belief that we developers hold about games and game players. We think we know what our audience likes because we know what we like. After all, we're not just game developers, we play games too, and we're convinced that this provides us with the insight to understand all players — and potential players — everywhere.

How many design discussions have you attended where somebody, in criticizing an idea, started a sentence with the words "Nobody cares about . . ."? If you've spent any time in the industry, the answer is probably in the dozens. Nobody cares about history (so the vast majority of games are set in science-fiction or fantasy worlds). Nobody cares about acting (so the acting in most games is abominable). Nobody cares about the story, everyone clicks through it (so the plot is trivial and the text is badly written). Girls don't play games (so very few games of interest to girls are produced). The list goes on and on. The basis for these sweeping statements is seldom any concrete evidence; it's just a belief that as game developers, we know what players want.

This logic is profoundly flawed. We may play games ourselves, but we are a peculiar class of gamers: those who also happen to be game developers. We don't represent those players who *don't* want to make computer games; there aren't any of them among us. We can't assume that our interests are the same as theirs.

Computer gaming is unique among entertainment media for the number of people that it inspires to want to make it their career. Most people watch TV without wanting to produce TV shows, and visit fairgrounds without wanting to run the carousel, yet a surprising number of people who play computer games also want to make them. Why should that be? It's because the games they're playing are designed for the kinds of people who are interested in making games — that is, they're designed by developers, for developers, or at least potential developers.

In my experience, market research in this industry is little more than a joke — and I used to work for one of the most successful publishers in the business. There are a few nods in that direction; every now and then somebody will collect up all the warranty cards returned by the purchasers and read what they've said, but any decent statistician would laugh out loud. Warranty card returns come from a tiny, self-selected minority of the customers. All they tell you is what the sort of person who returns warranty cards thinks — hardly a random sample. Oh, and we hold focus groups, but who do we invite to focus groups? Experienced gamers — specifically, the kinds of players who would enjoy spending an evening bending a publisher's ear. Again, not a terribly representative group. Other than that, the market research I've seen has been based on little more than hunches, conventional wisdom, and Myth #1.

Worse yet, there's another group of people we're ignoring entirely: the ones who don't play any games at all. Right now, we developers are all brutally clawing each other to sell more of the same kinds of games in the same limited shelf space to the same limited market of current game players. The real opportunity lies in selling to people who don't yet play games, but might start if they could find a game that they liked. These "proto-gamers" are the ones we should be reaching out to, the people we should be trying to create products for. But we don't know anything about them. All we know is they're not buying the games that we're making now — the kinds of games that *we* like.

There's a certain number of kids, mostly boys, who hang around the game store and buy a $40 game every few weeks. They're our traditional market. But there is a staggeringly huge number of people, mostly adults, and many of them women, who would like to take a few minutes between tasks at work or chores at home to play a light, fun, simple game that costs them a few cents, tops. Who's selling to them? Not most of us, that's for sure.

When I mention this notion to developers, especially young ones, I usually get a disgusted look and a flat dismissal: "I won't work on any game that I wouldn't want to play." The more thoughtful sometimes add, "I'm afraid I wouldn't do a very good job if I weren't passionate about the game." I sympathize with that notion — I, too, was one of those people who felt passionate about making games within five minutes of playing one for the first time — but ultimately it's short-sighted. It might be good art, but it's bad business. It still leaves you making games for game developers, and that's an overcrowded field. Passion's difficult to maintain when the game you spent 18 months building lasts three weeks on the store shelves.

Consider BARBIE FASHION DESIGNER. BARBIE FASHION DESIGNER was not constructed by ten-year-old girls. No offense to ten-year-old girls, but very few of them have the wherewithal to produce commercially viable entertainment software. BARBIE FASHION DESIGNER was developed by adults who were fairly unlikely to play with it much themselves. But in spite of that, they did an excellent job and they made a ton of money for themselves and their company. Not knowing the developers personally, I can only assume that they relied not on passion for playing the game itself, but on a different quality called professionalism, the desire to do a job well for its own sake. And it paid off in spades.

The way to overcome Myth #1 is to do something I have seldom seen done in the game industry: decide who your audience is upfront. Don't assume that you're selling to the same jaded old crowd and that you know exactly what they want. Instead, define your audience, then admit your ignorance about them. Go find some of them and actually ask what games they'd like to play, and where and how they want to play them. Who knows, you could discover a huge market that has been ignored for the last thirty years. Jackpot!

Myth #2: Realism Is Always a Primary Design Goal

What is a primary design goal? Here's how to find out. Make a list of everything you want to achieve with your game. These goals can be creative, technical, financial, anything. If you want your game to change the way the player thinks or feels, they

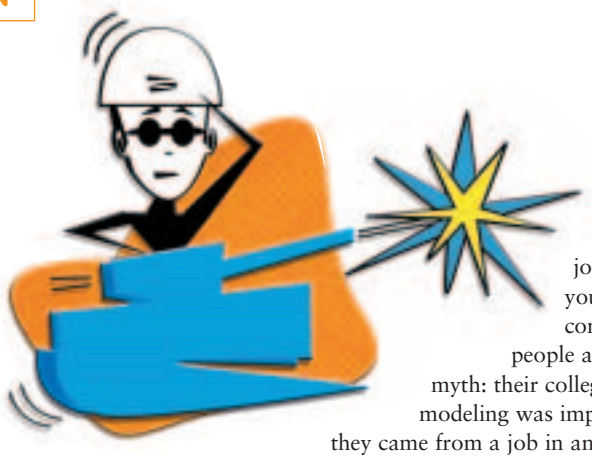

could be intellectual, emotional, or spiritual. They even could be political or social, if you want your game to change the world.

Now sort your list from the most important goals to the least. Then start from the top, run down the list, and ask yourself of each goal, "If we don't achieve this, will we consider the game to be a failure?" At some point, you'll stop saying "yes." Your attitude will change from "this goal is essential," to "this would be nice but isn't critical." At that point, draw a line across the list.

All the goals above the line you just drew are your primary design goals. They're the things you absolutely must achieve. The goals below the line are secondary, things you'd like to include but you don't feel you have to, and certainly not if doing so interferes with anything higher on the list. Secondary goals can be sacrificed for the sake of primary goals.

If you read enough game hype — advertising, box copy, press releases — you could quickly form the impression that realism is one of the most important features of any game. Publishers' marketing and PR departments are constantly harping on the subject, and if you pay attention to such things you might start to believe it yourself. Don't. It's Myth #2.

Realism can and should be a primary design goal in certain genres where it matters. High-end flight and racing simulators demand realism. Few other games need it, and there are some genres in which concentrating on it is actively detrimental. Many games are set in an artificial, symbolic world, and they should be. Monopoly has very little to do with the realities of buying real estate, and if you were to include those details (taxes, insurance, inspections, termites . . .) they would harm the game.

The primary source of this myth is pretty obvious: it's our own history. The audio, video, and processing capabilities of our machines have been continuously improving ever since the first computer game was written, and as a result, games are more realistic now than they ever have been in the past. More importantly, at every point in our history, the games have been more realistic than they ever were before. We're at the top of a steadily rising curve, and we always have been.

Of course these improvements look good, they sound good, and they serve to demonstrate technical competence and advancement. But they will occur automatically if your development team is taking advantage of the target hardware. Remember: primary design goals are those for which you are prepared to sacrifice something else, if necessary. If realism is a primary design goal, what are you willing to sacrifice for it?

Sometimes the answer is playability. Spectrum Holobyte's F-16 FALCON is an extremely realistic flight simulator, and when you play it in that mode, you find out why very few people are good enough to be fighter pilots. It's damned hard to play. Realism is F-16 FALCON's claim to fame, its raison d'être, so it's appropriate for realism to be a primary design goal, even at the expense of playability. But it's not appropriate for most other kinds of games.

Industry veterans probably don't need to be told this. The people who need to hear it are newcomers joining the industry from elsewhere, either young people fresh out of college, or people coming in from other industries. Those people are in fact the secondary source of this myth: their college professors taught them that accurate modeling was important in software engineering, and if they came from a job in another industry, it probably was. The guy who was brought in to head up EA's 3DO development team was a Ph.D. physicist, and he insisted that JOHN MADDEN FOOTBALL for the 3DO must have "realistic" physics. Unfortunately, until we changed his mind, this made the game unplayable. Because the athletes in a sports game are being guided by a simplistic game controller, you have to adjust the physics to compensate, but he didn't understand this. I was the designer of MADDEN at the time, and I taught the programmers a mantra to chant when they got in fights with their boss: "It doesn't have to be 'right,' it has to look good and play well." Eventually we won him over.

There's a legend from the early days at Atari that was told to me by someone who was there at the time. The arcade classic BATTLEZONE had come out, and the U.S. Army sent some people around to find out how Atari was making tank simulators for a few thousand dollars when they were paying millions for theirs. There was a meeting with a lot of brass hats on one side of the table, and a lot of long-haired, T-shirted, dope-smoking programmers on the other. The Army wanted to know how they achieved precision on such low-end gear. The programmers shrugged. "We don't," they said. The officers persisted. "So if an enemy shot should really miss your tank, but the computations are off and it hits it anyway . . . " "The player loses his quarter," the programmers said. "Big deal. He's not going to know, is he?" The Army decided to keep its own simulators.

Like most legends, the exact details may not be right but it's the message that matters, and it nicely illustrates my point about Myth #2. For the Army, realism was a primary design goal — a matter of life and death, in fact — and they sacrificed a lot of money to achieve it. For Atari, fun and manufacturing costs defined the primary goals, with realism a distant second. But BATTLEZONE was a great game for all that. This is an entertainment industry. Don't get needlessly bogged down with realism.

Myth #3: You Can Build a Hit by Imitating Another Hit

This myth tends to be believed more by business people than by designers. Designers usually want to innovate, not imitate. Still, if you look at the store shelves, there's a heck of a lot of imitation going on. It happens because a publisher's marketing and sales people notice that some competitor has had a hit, and they persuade the management that if the developers will just produce something like it, they can have a hit too. They're victims of Myth #3: the belief that you can build a hit by imitating another hit.

I'm sure we've all seen this myth at work. Someone will produce a brilliant new game, it'll be a massive hit at Christmas, and by the next E3 there'll be four or five schlocky knockoffs in the pipeline made by other people. They never look as good, because

chances are they're being rushed to market to catch the coattails of the original, and they never sell as well, because the "wow" factor is already gone. Why do they bother? Well, you can make a little money that way if you've got no pride and no creativity of your own. But in my experience the companies that do this are also-rans, second-rate outfits that will never really shine. For one thing, if the management has foisted it on the developers against their will, they are wasting their own talents. They have got their people building cheap knockoffs when they could be working on something innovative and new. No developer with any imagination is going to put up with that for long. They'll leave, and then the company has to find someone to replace them who doesn't mind working on cheap knockoffs. It's not a formula for building and maintaining an excellent staff.

Oddly enough, this can even happen within a single company. It occurs when the marketing department insists on creating a sequel to a game for which no sequel was intended. Sometimes a game is a hit because the development team has burned themselves out, put everything they had into that one game. If you then demand that they make another just like the first, you're bound to get an inferior product — they don't have anything left to give. With the current emphasis on franchises, we're usually better at product planning than that. But it still happens, especially with games that were unexpectedly successful.

I don't mean to suggest that you shouldn't study other people's games. I'm a firm believer in the value of studying other people's games; heck, I'm a firm believer in the value of studying everything. It's all grist for the creative mill. But there's a significant difference between keeping an eye on the competition and ripping them off. The latter is seldom successful. Our objective should be to surpass the competition, to create "wow" moments of our own, rather than hoping for a free ride on the back of someone else's imagination.

Myth #4: A Great Idea Will Make You a Fortune

Several times a year, I get letters from people who have great game ideas, but no clue how to make them a reality. They'd like me to teach them all about game development and marketing, but they're usually very cagey about what their idea is — they're afraid I might steal it and make a lot of money that's rightfully theirs. Alas, they've been seduced by Myth #4: the notion that a great idea will make them a fortune. It's a classic among fledgling game designers, so this section is for them.

Part of the reason people believe Myth #4 has to do with the way we're taught about the history of innovation. We're told neat little sound-bite chunks of history that don't include the whole story, things like "James Watt invented the steam engine." This gives the impression that in a world of horse-drawn coaches, James Watt saw the lid of his teakettle jiggling and suddenly the railroad was born. But James Watt was part of a much larger process, and what he really invented was a technical improvement to existing, stationary steam engines. He didn't invent the railroad, either. It was a man named George Stephenson who constructed the first steam locomotive — drawing on the work that Watt had already done, of course.

It's not that a great idea won't ever make someone a fortune. Sometimes one does, and a lot of people point to TETRIS as the perfect example. The problem is that truly great ideas on the order of TETRIS are extremely rare. For every TETRIS there are tens of thousands of seemingly great ideas that go nowhere. A winning lottery ticket will make you a fortune, too, but if you're serious about making money, lottery tickets are not the way to do it.

Another game that people point to as an example of a great idea that made a fortune is DOOM. Everyone remembers that DOOM was like nothing ever seen before, and it enabled John Carmack and John Romero to buy Ferraris. But DOOM wasn't actually a great idea from out of the blue; in fact, it's an excellent example of how fortunes are really made in the game industry.

The central idea of DOOM — running around and shooting things in the first person — was not new. There was a multiplayer game called MAZEWARS on the Xerox Alto word processor as far back as 1983, and there are probably earlier examples. The central display method in DOOM, a technique called raycasting, was not

new either. The reason DOOM did so well was not because it was great new idea, but because it was a brilliant execution of a variety of existing ideas. Raycasting may not have been new, but DOOM did it better than anyone else. The game was so simple, so fast and clean, that it still has its devotees today. I've seen it ported to all kinds of strange devices, even the LCD display of a digital camera.

That's the thing to know about fortunes in the game industry. There's nothing wrong with great ideas, but it's a mistake to think that they routinely lead to fortunes. What reliably makes money is not brilliant ideas, but quality workmanship: top-notch, first-rate, class-A execution. And the only way to obtain that is by the sweat of your brow.

Your chances of selling a great idea to a publisher just in idea form are rapidly vanishing. Once upon a time publishers might have bought ideas, but they don't anymore — and in any case, publishing companies are full of people who all have their own great ideas. Unless yours is of winning-lottery-ticket caliber, nobody's going to pay you for it. What publishers want is not ideas, but people who can create products, especially with the kind of quality that I was talking about.

But suppose you're still convinced that yours is a lottery-jackpot idea. How do you protect it? Here's a one-paragraph primer on intellectual property for wannabes. In America there are three ways to legally protect your work: copyright, trademark, and patent. A copyright protects a document of some kind, either text, a photograph, a sound recording, or some other expressive material. It doesn't protect the ideas in the document, only the document itself. You can't copyright an idea. A trademark is a name, slogan, or logo that you use to represent your company or its products. Trademarking something prevents other people in the same line of business from using it as well. You can't trademark an idea, though; again, it has to be something concrete. Finally, a patent is a means of protecting a new method for doing something. No one else must ever have done it before, and it actually has to be a method, not something abstract like a story or a character. You can patent an idea, but it has to be an idea about doing something, not just an image or a concept.

So if you have a game idea like "Robot Camels from Neptune Invade the Justice Department," you can't copyright, trademark, or patent it. You can draw a robot camel and copyright your drawing, and you can name the camels "Dromedroids" and trademark the name, but other than that you can't prevent other people from having and developing the same idea. Anybody else can make a game on the same subject. If you come up with some method of playing your game that has never been seen in any other game ever invented, you might be able to patent the method, but you still can't patent the robot camels. Besides, obtaining a patent is an expensive and time-consuming process. The burden of proof is on you to show that your method is so different from anything that anyone else is doing that you deserve exclusive rights to the idea. With game mechanics, that's going to be a tough sell.

The one other option is to treat the idea as a trade secret — that is, not to tell anybody about it, and if you do tell someone about it, to ask them to sign what's called a nondisclosure agreement, or NDA, first. An NDA is usually a simple, one-page contract in which somebody promises not to reveal your secret to anyone else, in exchange for getting a chance to look at it. Normally, no money changes hands. This is what I use when people want to consult me but don't want to tell me about their idea. I sign a nondisclosure agreement promising not to reveal their secret or use it for myself. However, getting signed NDAs doesn't take you any farther down the road to that hypothetical pot of gold, either.

All this may make it sound as if I think there's no point in innovation, and that we might as well go on making the same kinds of games because great ideas are worthless. Nothing could be farther from the truth. Great ideas are wonderful, but you need a realistic understanding of their value. No single idea is likely to be worth a fortune, so rather than clinging to it as if it were a diamond, we should continually generate new ones: learn, think, dream, create.

A couple of weeks ago I had the opportunity to visit a company called Hidden Dinosaur in Sweden. Its founder, Michael Stenmark, showed me his sketch book of ideas for the project they're working on. To say that I was amazed would be an understatement — overwhelmed is more like it. He had page after page of places, people, creatures, objects, stuff I had never seen the likes of. There wasn't an elf, samurai, or space marine in sight. Everything was fascinating and new, and every day he goes out and draws more things. He often travels for inspiration, and he never stops. He doesn't assume one idea is enough. They pour from his pen like a rainbow Niagara. He's got the right attitude about ideas: more is better.

If you've got a great idea, use it to practice your skills. Learn how to develop it. If you want to be a programmer, learn to program it; if you want to be an artist, learn to draw it; if you want to be a writer, write about it — and you can copyright all the material you create, so at least your labor is protected. You won't make a fortune, but if you work on it, your passion for it will show. Then you can bring it out at job interviews to demonstrate your talent. Don't worry about keeping it secret. In fact, do the exact opposite: talk about it, to anyone who will listen. If it's really that good, they'll be impressed with your imagination and more inclined to hire you.

In Conclusion

As I said, these four myths aren't lethal mutations. Believing them isn't necessarily going to destroy your career or your company. But like flaws in the genetic code, they accomplish nothing, they do more harm than good, and it's undesirable to pass them on to the next generation. By identifying and correcting them, perhaps we can effect a few repairs on our rapidly growing industry.