Is Justin Timberlake a Product of Cumulative Advantage?

By DUNCAN J. WATTS

As anyone who follows the business of culture is aware, the profits of cultural industries depend disproportionately on the occasional outsize success — a blockbuster movie, a best-selling book or a superstar artist — to offset the many investments that fail dismally. What may be less clear to casual observers is why professional editors, studio executives and talent managers, many of whom have a lifetime of experience in their businesses, are so bad at predicting which of their many potential projects will make it big. How could it be that industry executives rejected, passed over or even disparaged smash hits like "Star Wars," "Harry Potter" and the Beatles, even as many of their most confident bets turned out to be flops? It may be true, in other words, that "nobody knows anything," as the screenwriter William Goldman once said about Hollywood. But why? Of course, the experts may simply not be as smart as they would like us to believe. Recent research, however, suggests that reliable hit prediction is impossible no matter how much you know — a result that has implications not only for our understanding of best-seller lists but for business and politics as well.

Conventional marketing wisdom holds that predicting success in cultural markets is mostly a matter of anticipating the preferences of the millions of individual people who participate in them. From this common-sense observation, it follows that if the experts could only figure out what it was about, say, the music, songwriting and packaging of Norah Jones that appealed to so many fans, they ought to be able to replicate it at will. And indeed that's pretty much what they try to do. That they fail so frequently implies either that they aren't studying their own successes carefully enough or that they are not paying sufficiently close attention to the changing preferences of their audience.

The common-sense view, however, makes a big assumption: that when people make decisions about what they like, they do so independently of one another. But people almost never make decisions independently — in part because the world abounds with so many choices that we have little hope of ever finding what we want on our own; in part because we are never really sure what we want anyway; and in part because what we often want is not so much to experience the "best" of everything as it is to experience the same things as other people and thereby also experience the benefits of sharing.

There's nothing wrong with these tendencies. Ultimately, we're all social beings, and without one another to rely on, life would be not only intolerable but meaningless. Yet our mutual dependence has unexpected consequences, one of which is that if people do not make decisions independently — if even in part they like things because other people like them — then predicting hits is not only difficult but actually impossible, no matter how much you know about individual tastes.

The reason is that when people tend to like what other people like, differences in popularity are subject to what is called "cumulative advantage," or the "rich get richer" effect. This means that if one object happens to be slightly more popular than another at just the right point, it will tend to become more popular still. As a result, even tiny, random fluctuations can blow up, generating potentially enormous long-run differences among even indistinguishable competitors — a phenomenon that is similar in some ways to the famous "butterfly effect" from chaos theory. Thus, if history were to be somehow rerun many times, seemingly identical universes with the same set of competitors and the same overall

market tastes would quickly generate different winners: <u>Madonna</u> would have been popular in this world, but in some other version of history, she would be a nobody, and someone we have never heard of would be in her place.

Because it's not possible in the real world to test theories about events that never happened, most of what we know about cumulative advantage has been worked out using mathematical models and computer simulations — an approach that is often criticized for glossing over the richness of real human behavior. Fortunately, the explosive growth of the Internet has made it possible to study human activity in a controlled manner for thousands or even millions of people at the same time. Recently, my collaborators, Matthew Salganik and Peter Dodds, and I conducted just such a Web-based experiment. In our study, published last year in Science, more than 14,000 participants registered at our Web site, Music Lab (<u>www.musiclab.columbia.edu</u>), and were asked to listen to, rate and, if they chose, download songs by bands they had never heard of. Some of the participants saw only the names of the songs and bands, while others also saw how many times the songs had been downloaded by previous participants. This second group — in what we called the "social influence" condition — was further split into eight parallel "worlds" such that participants could see the prior downloads of people only in their own world. We didn't manipulate any of these rankings — all the artists in all the worlds started out identically, with zero downloads — but because the different worlds were kept separate, they subsequently evolved independently of one another.

This setup let us test the possibility of prediction in two very direct ways. First, if people know what they like regardless of what they think other people like, the most successful songs should draw about the same amount of the total market share in both the independent and social-influence conditions — that is, hits shouldn't be any bigger just because the people downloading them know what other people downloaded. And second, the very same songs — the "best" ones — should become hits in all social-influence worlds.

What we found, however, was exactly the opposite. In all the social-influence worlds, the most popular songs were much more popular (and the least popular songs were less popular) than in the independent condition. At the same time, however, the particular songs that became hits were different in different worlds, just as cumulative-advantage theory would predict. Introducing social influence into human decision making, in other words, didn't just make the hits bigger; it also made them more unpredictable.

So does a listener's own independent reaction to a song count for anything? In fact, intrinsic "quality," which we measured in terms of a song's popularity in the independent condition, did help to explain success in the social-influence condition. When we added up downloads across all eight social-influence worlds, "good" songs had higher market share, on average, than "bad" ones. But the impact of a listener's own reactions is easily overwhelmed by his or her reactions to others. The song "Lockdown," by 52metro, for example, ranked 26th out of 48 in quality; yet it was the No. 1 song in one social-influence world, and 40th in another. Overall, a song in the Top 5 in terms of quality had only a 50 percent chance of finishing in the Top 5 of success.

In our artificial market, therefore, social influence played as large a role in determining the market share of successful songs as differences in quality. It's a simple result to state, but it has a surprisingly deep consequence. Because the long-run success of a song depends so sensitively on the decisions of a few early-arriving individuals, whose choices are subsequently amplified and eventually locked in by the cumulative-advantage process, and because the particular individuals who play this important role are chosen randomly and may make different decisions from one moment to the next, the resulting unpredictably is inherent to the nature of the market. It cannot be eliminated either by accumulating more information — about people or songs — or by developing fancier prediction algorithms, any more than you can repeatedly roll sixes no matter how carefully you try to throw the die.

This, obviously, presents challenges for producers and publishers — but it also has a more general significance for our understanding of how cultural markets work. Even if you think most people are tasteless or ignorant, it's natural to believe that successful songs, movies, books and artists are somehow "better," at least in the democratic sense of a competitive market, than their unsuccessful counterparts, that Norah Jones and Madonna deserve to be as successful as they are if only because "that's what the market wanted." What our results suggest, however, is that because what people like depends on what they think other people like, what the market "wants" at any point in time can depend very sensitively on its own history: there is no sense in which it simply "reveals" what people wanted all along. In such a world, in fact, the question "Why did X succeed?" may not have any better answer than the one given by the publisher of Lynne Truss's surprise best seller, "Eats, Shoots & Leaves," who, when asked to explain its success, replied that "it sold well because lots of people bought it."

This lesson is not limited to cultural products either. Economists like Brian Arthur and Paul David have long argued that similar mechanisms affect the competition between technologies (like operating systems or fax machines) that display what are called "network effects," meaning that the attractiveness of a technology increases with the number of people using it. But even in markets that don't exhibit obvious network effects (like markets for low-carb or organically produced food, fuel-efficient vehicles or alternative energy technologies), sudden shifts in consumer demand can still arise, persist and then shift again. These shifts often come as surprises but are soon explained away as mere reflections of changing public sentiments. Yet while in some sense these markets do reflect what people want, that is true only of what they want right now. If markets not only reveal our preferences but also modify them, then the relation between what we want now and what we wanted before — or what we will want in the future — becomes deeply ambiguous.

Our desire to believe in an orderly universe leads us to interpret the uncertainty we feel about the future as nothing but a consequence of our current state of ignorance, to be dispelled by greater knowledge or better analysis. But even a modest amount of randomness can play havoc with our intuitions. Because it is always possible, after the fact, to come up with a story about why things worked out the way they did — that the first "Harry Potter" really was a brilliant book, even if the eight publishers who rejected it didn't know that at the time — our belief in determinism is rarely shaken, no matter how often we are surprised. But just because we now know that something happened doesn't imply that we could have known it was going to happen at the time, even in principle, because at the time, it wasn't necessarily going to happen at all.

That doesn't mean we should stop trying to anticipate the future, any more than we should stop trying to make sense of the past. But it does mean that we should treat both the predictions and the explanations we are served — whether about the next hit single, the next great company or even the next war — with the skepticism they deserve.

Duncan J. Watts is a professor of sociology at Columbia University and the author of "Six Degrees: The Science of a Connected Age."