

On the Research Value of Large Games

Natural Experiments in Norrath and Camelot

Edward Castronova
Indiana University

Games and Culture
Volume 1 Number 2
April 2006 163-186
© 2006 Sage Publications
10.1177/1555412006286686
<http://gac.sagepub.com>
hosted at
<http://online.sagepub.com>

Coordination game theory is quite important to a number of literatures but has had very few direct empirical tests because that would require experimental participation by large numbers of people. Large games however can produce natural experiments: situations that through no intent of the designer offer controlled variations on a phenomenon of theoretical interest. Unlike any other social science research technology, such games provide for both sufficient participation numbers and careful control of experimental conditions, making them like Petri dishes for social science. This article examines two examples of games as Petri dishes for macro-level coordination effects: (a) the location of markets inside EverQuest and (b) the selection of battlefields inside Dark Age of Camelot.

Keywords: *games; social science methods; natural experiments; MMORPGs*

Introduction

This article attempts to alert the reader to the research potential of large games, games with so many players and such rich complexity as to be hosts for genuine human social interactions. The claim is that such interactions are quite hard to study at the scale of all society. Until now, it has not been possible to take all of society as a research object; such a thing is too big to fit in a lab, and the histories it produces are never replicated anywhere. Thus, although we might believe theoretically, historically, and ethnographically that society operates a certain way, and we might have small-scale experiments that support our beliefs, it has generally not been possible to observe whole societies under controlled conditions. Now however, with the advent of synthetic world technology, it is indeed possible to replicate entire societies and allow them to operate in parallel. This article is an exploration of what might be learned from a research program built around such a method. Having no lab (as yet) in which to conduct this kind of experiment, we turn instead to the

Author's Note: Please address correspondence to Edward Castronova, PhD, Associate Professor and Director of Graduate Studies, Telecommunications Department, Indiana University, 1229 E. 7th Street, Bloomington, IN 47405-5501; e-mail: castro@indiana.edu.

synthetic worlds that already exist and ask whether any of them have produced quasi-experiments: situations in which important factors have been controlled unwittingly by the designers. Two specific examples of such quasi-experimental circumstances will be put forth, one in EverQuest (EQ), by Sony Online Entertainment, and the other in Dark Age of Camelot (DAoC), by Mythic Entertainment. Both cases involve the theory of coordination games. In what may be the first finding of its kind, these cases provide controlled-experimental evidence on the level of all society that coordination effects do exert force there. In other words, coordination effects are not just of theoretical importance, and our evidence of them from ethnography, history, and small-scale experiments is not misguided. They are genuinely powerful at the most aggregated level of human society. At the same time, the ease and clarity with which they could be observed speaks to the immense value of large games as a research tool, prompting perhaps a rethinking of how social scientists do their business.

The article has six remaining sections. First I introduce the concept of coordination games and their theorized role in social outcomes. The next section explains why any current evidence about coordination effects at the macro-social level has to be indirect and thus fairly unpersuasive. The next section puts forth the idea of large games as a new research tool that allows direct, replicable observation of whole societies in operation. The following two sections present evidence of coordination by players in the large games EverQuest and Camelot, respectively. The final section concludes with a call for a new methodological direction in social research: the use of large games to study behaviors of the whole.

The Power and Fragility of Coordination

You and a friend go to a large shopping mall but you become separated in the crowds. Neither of you has a phone or any other communication device, and the mall does not have a public address system by which you might page one another. To meet your friend again, where would you go? Or would you go anywhere at all? Perhaps you should stay put and wait for your friend to find you. But what if she does the same thing? On the other hand, if you both decide to go roaming, she will probably pass through the spot you're standing in 10 times before you accidentally run into each other. Perhaps you should think about the most likely place she would look and go there. Maybe the last store she was in. Or the last store that you were in. Or maybe the exit or even the car is the best bet. Maybe the mall has a "Meeting Point," a spot created just to minimize the costs of this dilemma. Suppose you looked at the directory and there was a red dot labeled *LOST PEOPLE MEET UP HERE*. Well, you could probably bet that she would go there, and therefore, so should you. Furthermore, if she is following the same reasoning, she's going to place a high probability on the red dot too, which will make it her first choice. And—coming

back to you now—these guesses of yours about her reasoning further raise the likelihood that the red dot is the place to go. And so: You go to the red dot. And there she is.

The red dot, the Meeting Point, is an example of a coordination equilibrium, a social phenomenon that shares two properties that social phenomena don't often share: power and fragility. Power comes from the fact that human coordination saves lots of time and money, and it gets remembered. The next time you and your friend go to the mall, you'll know where you'll meet up if you get lost again: the Red Dot. But if the mall did not have a Red Dot, and you had ended up meeting at, say, the North Exit, then in all likelihood you'd always and forever find one another at the North Exit. Coordination equilibria, also known as focal points, place a structure on our expectations of what others will do, and those expectations of others have an incredibly powerful effect on our choices. I work at Indiana University, and I wear red clothes to the football games. Is it because I like red clothes? No. Is it because the chancellor will reduce my pay if I wear blue? No. Is it because I want to march in step with my fellow Hoosiers, in perfectly jackbooted style? No. It's because I have observed rude behavior by sports fans in reaction to nonconforming sports attire (all in good fun, but still, it's uncomfortable), and my understanding is that everyone at this particular stadium will be wearing red. If I went to a stadium in Ann Arbor, I'd wear blue. In East Lansing, green. In Evanston, purple. It is the culture of sport to coordinate on team colors. And for each team, the color is the focal point, a coordination equilibrium that exerts incredible pressure on hundreds of thousands of people, forcing almost all of them to wear some arbitrary, even ridiculous color (purple?) whether they really want to or not.

There's a second aspect of power worth mentioning. Someone at Indiana University, back in the mists of time, decided that our colors would be red. If you happened to be a red dye maker in Bloomington, that was a happy choice with immense implications for your business. At the shopping mall, someone decides where the Meeting Point will be. If it happens to be right in front of your store, then that decision maker has just directed some traffic to your business. Not every coordination equilibrium happens intentionally, but when someone does choose the focal point, that person exerts influence on human behavior, directing it into specific channels.

And yet this second source of coordination power is also the source of its fragility. In many social situations, there are many conceivable foci. Take colors. You could have red, blue, yellow, green, purple, black, orange, or any of an infinite number of combinations. If one is selected—red—then all people coordinate to red. Yet the moment that another is selected and becomes known as the new focal point, the first loses its power completely and utterly. When the St. John's University basketball team changed their mascot from Red Men to Red Storm, anyone whose business relied on the Native American imagery lost something, whereas those who made T-shirts with clouds on them gained something. And this can be an instantaneous effect; coordination systems can cause huge flows of well-being in mere seconds.

Examples other than sports and shopping serve to illustrate the power and fragility of coordination at a somewhat more meaningful level. Governmental power for one is a coordination equilibrium. In times of revolution, all of society must determine whom all of society will consider “The Government.” That is why the putschists seize the radio station—by broadcasting to the entire country that Herman is the new dictator, they instill the expectation that everyone is going to be treating Herman as the new dictator. And if everyone treats Herman as the new dictator, well, he is the new dictator. His orders will be followed because everyone expects that those who do not follow his orders will be killed. Thus, heads will roll if Herman pulls off this shift in the focal point. His head will roll if he does not. It is often difficult to understand how entire systems of absolute power, such as the Ancien Regime in 18th-century France, could be toppled in the course of a few days, but in fact, it’s merely the switch from one governing focal point to another. And simple, even localized events, such as the storming of the Bastille or the opening of the Berlin Wall, are sufficient to signal to the entire world that the focal point has changed. Once everyone believes that parliament and not the king has ultimate authority, the parliament, by virtue of those shared beliefs, does have ultimate authority. The king becomes a mere citizen indeed, subject to the soldiers’ orders and not vice versa.

The value of money is another coordination equilibrium. The dollars in our pockets have value because we all expect everyone else to treat them as though they have value. As I write this, the focal point is that one U.S. dollar is worth about one bottle of Coca Cola. That is: I expect that everyone will treat my dollar as though it has the same market value as a bottle of Coke. Factually and objectively, the dollar has no value whatsoever—you can’t eat it, you can’t write on it, you can’t wear it. It’s worthless in use. But in exchange, it is worth something, and what it is worth is determined by the coordination equilibrium. If I decided that my dollars were worth two bottles of Coke, not just one, I would be off the equilibrium and would not succeed in any attempt to get people to trade with me. If I decided my dollars were worthless and gave them away for nothing, I would again be off the equilibrium and would find myself dead of starvation fairly quickly. Therefore I (and by extension everyone else) choose to act as though the dollar is worth one Coke, and I can expect that everyone else will as well. From this perspective, inflation can be seen as the gradual shift in the focal point, a kind of massive dance in which all of society unconsciously takes the same step of -3.4% in its understanding of what a dollar is worth. This dance is ultimately being done to a rhythm called out by the government of course. And when the government does a poor job as bandleader, the dance steps get trickier, larger, and the dance of inflation becomes a terrifying and conscious act, sometimes with dramatic conclusions, such as during the German hyperinflation of 1922-1923. Wages were posted not annually but daily, even hourly. A new coordination game erupted around not the value of the money but on the daily percentage decline in the value of the money: “How much are we stepping today? I need to align my product prices with my workers’ wages!” A government can beat the drum so

quickly that all pensions are wiped out or that all debts become meaningless, being repayable in a currency worth only a tiny fraction of what it was worth when the debt was incurred. The value of money, like governmental authority and all coordination equilibria, is both powerful and fragile.

A third example of the importance of coordination is in media. Take journalism. Once one powerful outlet decides that a story is important, all outlets consider the story important, and they therefore run with it. Thus, from the standpoint of the viewer, it is important. Woe to the outlet that fails to run an important story or that runs stories that no other outlets think are important. This then drives all the outlets to guess what all the outlets will think is important, a coordination game at the metalevel that is only loosely connected to what individual people think is important. And because we as individuals have a natural desire to seem “on top of things” to our peers, it is not at all clear that our own expressions of what is important reflect any internal judgment of relevance. How relevant is it to any one person whether Tom Cruise or Tom Hanks wins the Oscar for best male performer? Yet almost everyone has an opinion, has thought about it. Such a thing is important because it is important. With coordination phenomena, what seems to be circular logic is in fact all the logic there is. It is, if you will, a brittle unlogic and yet quite powerful.

The aforementioned examples illustrate that coordination games are an important element of thinking about economics, politics, and culture. We know a great deal about them analytically because of the work of mathematical game theorists. This approach, defined largely by John von Neumann, Oskar Morgenstern, and John Nash, favors formal mathematical expressions of players, strategies, and outcomes and assumes that players will be rational in their choices of what to do. Perhaps the best overall analysis of coordination games was written by recent Nobel laureate Thomas Schelling (1960): *The Strategy of Conflict*.

From Schelling (1960), a coordination game can be expressed as follows. There are two players, 1 and 2. Each player chooses either move X or move Y. If both choose X, Player 1 receives M as a reward, and Player 2 receives M + 1 (M could be dollars or power or puppies—anything the players believe is a good thing). If both choose Y, Player 1 receives M + 1, and Player 2 receives just M. If one chooses X and the other chooses Y, they both receive nothing. This game can be expressed in a choice diagram in Figure 1.

This simple game exhibits all the essential features of the coordination phenomenon. First, whatever is the thing enumerated by the payoff number M, we have assumed that it is something the players like. Having some of it is better than having none of it. Thus, both players would rather have a coordinated outcome—[X, X] or [Y, Y] rather than a mix of the two—because only when they choose the same move do they get anything positive. Yet what is unclear is which of the two coordination points they will choose. And that also matters. If they get together on [X, X], Player 2 is going to be slightly better off, getting one more of the good thing. If instead they get together on [Y, Y], that’s going to be better for Player 1. Just what

Figure 1
A Coordination Game

		Player 2	
		Move X	Move Y
Player 1	Move X	Player 1: M Player 2: M + 1	0 for both
	Move Y	0 for both	Player 1: M + 1 Player 2: M

they will do depends on other things that determine how the game gets played. For example, suppose Player 1 gets to move first. Her best bet would be to choose Y. If she chooses Y, Player 2 then has the choice of X (receiving nothing) or Y (receiving M). Because M is better than nothing, Player 2 chooses Y, thus giving Player 1 M + 1. You might say that Player 1 “wins” this one. If Player 2 was given the first move, he would “win” because his first move would be X, forcing Player 1 into X as well.

Another possibility for the structure of the game puts in the hands of a third person. Suppose the two players cannot talk to one another and they both have to move at the same time. This would be rather like the shopping mall example at the start of the article. Not having any idea how the other player will move, both players will make some essentially random choice, winding up with one of the four outcomes with about equal probability. Now suppose some third person can talk to both of them. That third person might say “Make Move X,” “Make Move Y,” or “Make a random move.” Suppose the third person says “X.” In that case, both players should reasonably conclude that the other player is more likely to do X than anything else. Therefore, X makes more sense as their own move. And by this same reasoning, X is seen as a still better move for the other player. And once more, that makes X yet better again as their own move. The circularity of reasoning reduces to the difficult truth that X is chosen because X is chosen. But recall that the system moved to X and not Y only because of the third person’s statement. This is the analytics behind the Meeting Point and the power of its Big Red Dot to make two friends meet at a certain point in the mall. This also illustrates the power of media to induce behavior and the power of a dictator to normalize his own power. Through properly constructed and positioned broadcasts, those who own microphones can make the whole world expect a certain behavior, and because that behavior is expected, it will be performed.

Another structure for this game illustrates the power of a history. Suppose Player 1 and Player 2 are members of two cultures, Culture 1 and Culture 2, that have been having this kind of interaction for a very long time. Somewhere back in the mists of

the past, at a time that no one remembers, it became established that people from both cultures would play Y, giving a slight advantage to members of Culture 1. Every once in a while, a Culture 2 member might try to upend the apple cart and play X instead, but in those cases that person always came away empty-handed because the Culture 1 people always play Y. X and Y give everybody nothing, making it seem fruitless for Culture 2 to try to change things. And yet, those who truly understand this interaction cannot avoid realizing that it could very easily go the other way. But it does not. The heavy hand of history forces Culture 2, bitterly, to accept a disadvantaged outcome. The only way to change things: a revolution in expectations that starts a new history. Only if everyone in Culture 2 and Culture 1 can be made to expect the X move will all the players begin making X moves. On the other hand, all one has to do is change the expectations and nothing more. If you can change everyone's mind about what the "normal" move is, suddenly X becomes the normal move, and Culture 2 becomes the privileged culture.

Coordination games are thus fairly simple to understand analytically, and armchair empiricism suggests that they are prevalent in a number of important social phenomena. In terms of their influence over human affairs, they seem to be both powerful and fragile at the same time. A coordination equilibrium, or focal point, is immensely powerful so long as everyone believes in it. Yet a slight alteration in these beliefs is all it takes to destroy the power of a focal point quite completely. A system with these features creates a very unique pattern of hopes and fears in anyone trying to work within it. Changing the collectivity of expectations is in one sense impossible, in another sense easy. It seems impossible because people are autonomous—who can change the mind of every person simultaneously? And yet it seems easy because change does not require one to move mountains or harness the sun—you just have to make people have different expectations about one another's behavior.

Empirical Evidence?

It would certainly be broadly useful if we knew that this fragile yet powerful piece of coordination unlogic was in fact the reigning dynamic behind real-world social change. The theory is fairly easy to disbelieve; it defies intuition at a number of levels. I've run through the model with my wife in an attempt to coordinate on a simple plan in the event we get separated in a crowd: She should stay in one spot while I roam around. Or was it, I should stay in one spot while she moves around? When we get lost, you see, I invariably forget what the coordination actually was. And so we roam around more or less aimlessly like all couples do. In other words, it's not clear that real people are as rational and mindful as the theory requires. When people are involved in something that looks roughly like a coordination game, do enough of them choose behaviors that are "rational" so that the system evolves as the coordination theory predicts?

It may be hard to believe, but even after five decades of work on this subject, there's not much solid evidence that coordination game dynamics happen in the real world.

How can this be so?

The core problem lies in the nature of social science as a research field. Social scientists study humanity in its groups. There are really only a few ways to do this.

1. You can write abstract theories using words or mathematical symbols and "test" them against the intuition of the reader. Readers are all members of human societies. If the assumptions make sense to them, and the logic looks tight, the argument can be accepted.¹
2. You can examine real human societies and compare across time periods or geographic boundaries. If two societies seem roughly similar in most respects, the remaining differences may relate causally to one another.²
3. You can conduct experiments with groups of people, assigning them at random to experimental and control groups. You expose the experimentals to some intervention, making that intervention the cause of any subsequent differences between the groups.³
4. You can immerse yourself in the society in question and use your deep knowledge to render subjective opinions about what is happening.⁴

All of these methods run into severe problems when asked to help us understand large-scale social phenomena. Ethnography (4) suffers from a lack of a sufficiently broad perspective for the study of billion-person societies. Pure theory (1) forces too much abstraction from the incredibly complex reality of a billion-person world. Comparison (2) cannot find two societies similar enough to allow observed differences to be traced to any set of causes capable of description or summary. Experiments (3) do not have enough participants. All of these methods can be used to hint at, point to, and suggest something about large-scale social entities, but none of them deals effectively with large-scale social entities themselves.

Today, we study large-scale social entities only indirectly. There are no laboratories in which large-scale social entities can be directly examined. A social scientist today is like a marine biologist, an expert on salmon, who can examine anything in the world—except salmon. He can theorize about salmon, he can watch movies of historical salmon that are now dead, he can conduct experiments with minnows and crabs. He can go out and swim upstream with salmon. But what he cannot do is bring some salmon into his lab or put ID tags on them in the wild. He can only swim upstream with them and hope he's got an accurate theory about why they are doing that.

Large Games: Petri Dishes for Social Science

In recent years, a new technology has emerged that represents the social science equivalent of a petri dish, or a supercollider: an expensive machine that provides the only way to directly study certain interesting atomic phenomena. If you want to study the

properties of atoms as they bang together, you must either do it indirectly or build a big machine that can bang atoms together under controlled conditions. Similarly, if you want to study phenomena operating at the level of whole societies, such as large-scale coordination games, you must either do it indirectly or build a big machine that can host large-scale coordination games under controlled conditions. Large games do exactly that.

What is a *large game*? I propose it as a technical term to replace a rather ugly piece of jargon, MMORPG, which stands for massively multiplayer online role-playing game. MMORPGs are a subset of the universe of synthetic worlds that the entertainment industry has started to build in great numbers. Synthetic worlds are online spaces that many people can use at once. If their use is channeled through a character that operates in a constructed social environment, complete with markets, governments, transportation, armies, and so on, we say that the synthetic world is actually a kind of game. *Game* here is understood in the sense of Malaby (2005), a designed choice environment with an uncertain outcome. And yet because this game occupies hundreds of hours of time of hundreds of thousands of people, we should say that it is not an ordinary game, it is a large game.

The defining feature of a large game is that because of its sheer size and complexity, it can be categorized as a genuine human society. In other words, the society we all lived in before the advent of synthetic worlds was itself a large game. True, it was the only large game, but these days it is no longer unique in that sense. The large game of Earth society now competes for my time with the large game of Norrath society, Norrath being the world of the video game EverQuest, which at this writing hosts the attention of some 400,000 human minds on a daily basis. And so long as we focus on the core mechanics of a large game, there will be little in the way of significant difference between the behavior of individuals in one large game or another. In other words, core phenomena such as supply and demand, patronage, stigma, diplomacy, and yes, coordination ought to operate similarly in all large games wherever they are located. When teams of engineers from different companies tackle a design problem involving automotive emissions, the people are real, they are really engineers, the car is real, and the pattern of disjointed, competitive teamwork is real. When groups of warriors attack families of dragons in Norrath, the people are real, but their status as “warriors,” arguably, is not. Neither are the dragons “real” in some sense.⁵ But the pattern of disjointed competitive teamwork is real and no different in its essence than the teamwork of the Earth engineers. Competitive, overlapping teamwork is just one of an uncountable number of core phenomena shared by all human societies, and if a game gets big enough, these phenomena will appear in them with absolute fidelity.

We see the fidelity of large games to the large game of Earth society for example in the market forces that have been the subject of much of my earlier writing (Castronova, 2001). It is indisputable that supply and demand operate in synthetic worlds like EverQuest in exactly the same way they do in San Diego. If one has any doubt at all about this, the questions to ask are the following: Did the invisible hand

operate in ancient Athens? Does it function the same way in contemporary Nairobi? What about in Costa Rica, Beijing, or London? Of course. And, is it possible to imagine a more widely disparate set of cultures than these? No, not really. Of course the invisible hand operates in all these environments. And therefore of course it operates the same way inside all large games. If it did not, what good would it be as a theory of human interaction? Any good theory ought to function in about the same way everywhere.⁶ Therefore any theory worth examining ought to function about the same way in all large games, online or offline. And there is no doubt in my mind that the core theories of economics all work the same way whether the societies involved have collected themselves somewhere on Earth or somewhere in cyberspace.

What's different about the online societies is that the large games being built out there operate under controllable conditions. This is indeed a dramatically important difference, whose meaning for the future of social science research can hardly be overstated. Large games are social science petri dishes. They are synthetic worlds within which real societies accumulate. And although large games are certainly expensive to build, those who do build them manipulate many thousands of levers, buttons, and knobs that affect everything about the way the resulting society operates. Prices can be high or low, communication can be easy or hard, travel may be rapid or slow. There might be voting, or there might not be. It can rain or shine. Builders can vary almost everything except human nature itself, making users mediate their interactions with one another through almost any kind of environment imaginable.

Given this level of control, an easy and yet breathtakingly powerful research strategy almost immediately leaps to mind: Build several synthetic worlds in exactly the same way except for some difference in a variable of interest. Put hooks in the data stream so that relevant metrics within the container societies can be observed. Attract people into the worlds, sit back, and watch what happens. With controlled construction of synthetic worlds, what we have in effect is a laboratory for large-scale social science research.

Although a research method of controlled experiments in large games makes sense in theory, we do not know if it would be feasible or useful in practice. One might ask whether there is any evidence that large games offer enough of a controlled environment to make this approach really useful. For example, is it possible to introduce enough variance in the experimental effect to induce significant and observable behavioral changes? Perhaps behavior is so dominated by the surface-level "game," the whacking of orcs and so on, that the things we care about (markets, politics, etc.) are mere sidelights that never change much no matter what one does.

Thus, we have two related questions to examine:

1. Do coordination effects really happen on a large scale in human society?
2. Are large games really amenable to controlled social experimentation?

The astute reader senses that the evidence I'll be submitting favors answers of "YES" and "YES."

Evidence From Norrath

In August 2001, I conducted a survey of EQ players. The Norrath Economic Survey, or NES, became the basis for an initial article on that economy's core statistics (Castronova, 2001). One of the questions I asked, on which nothing has been reported thus far, involved the location of the player-to-player market. In EQ at that time (there have since been revisions; see following), trade happened as follows. A seller would announce in the auction chat channel what items he had for sale. Because EQ's terrain was split into dozens of separate zones, this announcement could only be "heard" (that is, seen and read in the chat box) by those in the same zone as the seller. At the same time, everyone in that zone would hear it. A potential buyer would hear the proposed sale and respond by private message to the seller. The two would then negotiate privately, striking a price. Then they would meet—that is, place their avatars in close proximity to one another—to enable to exchange functionality of the interface. One performed exchanges by clicking on another person's avatar, which opened a trade window. The buyer would place money in his side of the trade window, and the seller would place the item in his side. Both sides could examine the goods being proposed for exchange. When satisfied, buyer and seller would separately click an Agree button. When both Agree buttons had been clicked, the items would change possession. The purpose of my question about market location in the NES was simple: Given this rather cumbersome means of market making, were there certain zones that were preferred by all sellers and all buyers? Did the player market as a whole have a specific geographic location?

The question of location bears interest because under the conditions of trade in Norrath, EQ's territory, the location of a player market looks like a coordination game. Consider the following unlogic: "The market is here because the market is here." Once the market exists in location X, it clearly behooves all players to travel to that location to conduct market activities. Buyers will come to X to shop because that's where the sellers are. If you want to hear all the offers and take advantage of seller versus seller price competition, you should shop in Zone X. Moreover, if you shop in X, doing the transfer will be quite easy because both buyer and seller avatars are already in the same zone. And because the buyers go to Zone X to buy, sellers will go to Zone X to sell. That's where the buyers are. For reasons of both information exchange and transactions costs, it behooves all market participants to do their market activity in the same spot. And if everyone expects that other market participants will get together in Zone X, then everyone will indeed get together in Zone X, confirming the expectation and coronating Zone X as the marketplace. And it is important to stress that nothing in this theory of market location is specific to online role-playing games. It is equally amenable to the large game of Earth. If coordination has a role to play on Earth, it has one to play in Norrath.

The question about market location was structured so as to reveal commonly held expectations about where one should sell things. For reference, here is the question as it appeared in the survey:

14.A. If your main character had a somewhat uncommon item that you needed to sell quickly at a fair price, where would you go to sell the item? What would be the best zone? Assume that you could get a teleport directly to that zone, so that travel time does not matter. Indicate a single zone here: _____.

In essence, the question asked respondents for the preferred place to sell an item independent of where their main character happened to be at the time of the survey. The answer was an open field; respondents could have put down anything.

It is interesting therefore that the responses came in the pattern shown in Table 1: Server by server, EQ players agree—to a statistically extreme degree—where the main market is.

It is unusual in any distribution of human responses and behaviors to find commonalities above 65%. The populations of EQ servers usually ran to 10,000 players, of which some 2,000 to 3,000 would be online at any one time.⁷ When more than 85% of such a population says that they share some feature, that feature is no longer just “frequent” in that population; from an outside perspective, we should think of it as “near universal.” That feature certainly dominates culture and expectations. Thus, the fact that a large majority of EQ’s server populations have 85% to 95% agreement on the location of the market—in response to an open-ended question, no less—is strong evidence that the zones indicated represent a powerful shared expectation of these populations. “The market is there.” Such statistics bear up my own impression when wandering the lands of Norrath: On every server, everyone knew where the market was. It was a firm and fixed part of the culturo/economic landscape. Where the statistics seem to offer less certainty about market location, there are some clear explanations. In the case of Lanys T’Vyl (36% agreement) and Vazaelle (65%), the response rate was uncharacteristically low (14 and 57 observations, respectively). And in the case of Tallon Zek and Vallon Zek (69%), we have here the hard-core player versus player (PvP) servers. It is a prediction of market theory that conditions of war impede efficient market formation. From a coordination game perspective, if coordination leads to vulnerability to attack, as it would on these servers, this danger lessens the incentive to coordinate. Better to throw an item away than earn death by trying to sell it. As a result, the forces directing market activity into a single location are lessened; the marketplace does not form as strongly as it does under conditions of peace. This also is consistent with the general idea of market location as a coordination game. Thus, whether one looks at the exceptions or the consistent cases, the rule is quite clear: In Norrath, market location looks like a coordination game.

Now, in what sense is this kind of evidence more persuasive than, say, historical evidence from Earth? For we can find many examples of markets being located on Earth in certain places. America’s principal stock exchange is located in New York and has been for decades. Chicago is the location of America’s primary commodity market. The television industry is located in Los Angeles. The question is, Does this kind of evidence identify coordination effects as the source of the staying power of these market locations? Having a market in a single place is certainly consistent with

Table 1
Distribution of Market Locations by Server

ID	Name	Launch Date	Split to	Split From	Unique Features	Market Zone	Market Zone Frequency	N
01	Bertoxulous	01/04/99	34			EC	84.78	92
02	Fennin Ro	02/17/99	27			EC	89.60	125
03	Rallos Zek	02/17/99			PvP	EC	89.32	103
04	Rathe	02/17/99	39			EC	92.04	113
05	Tarew Marr	02/17/99	37			EC	93.46	107
06	Xegony	02/17/99				NF	92.31	91
07	Cazic Thule	03/03/99	38			EC	89.61	77
08	E'ci	03/03/99	39			EC	90.00	80
09	Karana	03/03/99	38,34		Europe	EC	95.96	99
10	Mithaniel Marr	03/03/99	36			EC	91.74	109
11	Povar	03/03/99	37			EC	83.33	96
12	Solusek Ro	03/03/99	38			EC	89.04	73
13	Veeshan	03/17/99	30			EC	90.77	65
14	Bristlebane	03/26/99	30			EC	87.63	97
15	Innoruuk	03/26/99				NF	88.30	94
16	Tunare	03/30/99	36			EC	92.05	88
17	Rodcet Nife	04/28/99				GF	95.89	73
18	Prexus	06/07/99				EC	86.27	102
19	Brell Serillis	07/21/99				EC	91.54	130
20	Tallon Zek	07/30/99			PvP	GF	68.89	45
21	Vallon Zek	08/02/99			PvP	EC	69.05	84
22	Tribunal	08/17/99				NF	90.12	81
23	Nameless	10/14/99				GF	81.71	82
24	Lanys T'vyl	12/21/99				GF	35.71	14
25	Vazaelle	12/22/99				GF	64.91	57
26	Luclin	01/23/00				EC	85.92	71
27	Druzzil Ro	03/01/00		2		EC	91.67	84
28	Morelle-Thule	04/18/00				EC	90.32	93
29	Torvonnilus	04/18/00				EC	85.45	55
30	Saryrn	04/28/00		14, 13		EC	93.33	105
31	Seventh Hammer	04/28/00				EC	85.26	95
32	Erollisi Marr	05/06/00				EC	90.16	61
33	Quellious	06/24/00				EC	83.61	61
34	Terris-Thule	07/19/00		09, 01		GF	91.58	95
35	Tholuxe Paells	09/28/00				GF	68.75	64
36	Drinal	10/03/00		10, 16		EC	88.37	86
37	Xev	11/21/00		5, 11		NF	91.43	70
38	Ayonaе Ro	12/19/00		07, 09, 12	Europe	GF	94.00	100
39	Zebuxoruk	03/21/01		08, 04		GF	80.00	85
40	Sullon Zek	07/02/01			PvP	EC	93.15	73

Source: Castronova (2001).

Note: Splits refer to the practice of founding new servers by allowing players from certain older, crowded servers to move to the new server. Example: On March 1, 2000, the server Druzzil Ro was founded, and all players from Fennin Ro were allowed to transfer characters there. PvP = player versus player. Market zone: EC = East Commons, GF = Greater Faydark, NF = North Freeport. Market zone frequency = percentage indicating the given zone. N = number of respondents from that server.

coordination game theory, but it's consistent with many other theories too. One could tell stories. L.A. has TV because L.A. had movies, and L.A. had movies because it had sun. But San Diego has sun too. However, San Diego did not have the Southern Pacific (SP) Railroad, and that is because L.A.'s business leaders convinced SP to make their port and not the (superior) Port of San Diego the railroad's western terminus. Was this because L.A. was closer to San Francisco? If so, then L.A.'s status as the headquarters of television was driven by the forces of economic history, not coordination effects per se. Or was it an accident, a random choice in a coordination game being played by lots of folks in the late 19th century, among them the directors of Southern Pacific? In the latter case, if history were to repeat itself, coordination effects alone would allow a TV industry to emerge in La Jolla instead of Burbank. Similarly, if we merely note that for months and months people on the Bertoxulous server have been doing their market activity in East Commons (EC), we have added very little to what we know about coordination effects.

Where EQ offers some new information on coordination is in the conditions under which its marketplaces were founded: In EQ, each server is an exact replica of the others, and so the EQ servers as a whole represent 40 replays of the same history. In other words, on Earth we have one experimental world, the Earth itself. In Norrath, we have 40 Earths, each exactly the same except for a few differences involving age, the rule set, and the uncontrolled sorting of populations to servers. Thus, when we observe a pattern of variation by server, we are observing the grand and exhilarating WHAT IF that has been dreamed of by generations of social observers: What if we reran the history of television? Would it end up in Burbank or La Jolla?⁸ With 40 Norraths, exactly that kind of judgment is possible; we can do more than tell tales based on a single history, we can observe multiple, parallel histories and weave an argument based on a firm understanding of underlying probabilities. And this is only possible because EQ is not just a large game, it is a large game occurring under controlled circumstances of parallelism. Controlled to emerge in 40 replicated incarnations, the worlds of Norrath offer unprecedented new information about social behavior.

The specific information offered about coordination is quite helpful and clear, in fact: The evidence from Norrath suggests strongly that coordination plays a very strong role in market location. This is because we can view the evidence from Norrath as the outcome of an experiment. The experiment is this: Open 40 versions of the same physical world and attract individuals into them. Observe where markets are located. If market location has the features of a coordination game, as illustrated in Figure 1 earlier, it should be the case that

- the market coalesces at a single location, AND
- the market does not always coalesce in the same place.

The first prediction is consistent with the observed history of our offline large game, the Earth economy. L.A. is a single place, and the TV industry is there, not elsewhere.

The first prediction also is confirmed in our Norrath large games: On each server, the market has a single location.

The second prediction is where EQ provides a unique perspective, allowing us to identify the distinct role of coordination effects. For if market location is about coordination to some degree, then there is the possibility of multiple, different outcomes. Indeed, it is possible for different market locations to become entrenched with an equal degree of force even though the locations are better for some people and worse for others. And indeed, that is the pattern we observe in Norrath.

Returning to Table 1: On most servers, 27 to be exact, the market is located in East Commons. On 9 servers, the market is in Greater Faydark (GF), and on the remaining 4 it is in North Freeport (NF). We who know EQ well understand the various advantages and disadvantages of these locations as marketplaces.⁹ Arguably, EC had the most advantages, but in the absence of parallel experiments, we would not know whether the advantages of EC were such that coordination effects could not displace EC as the marketplace. However, because we do have parallel experiments, we can conclude with certainty that mere coordination was sufficient in 13 of 40 cases to keep the market elsewhere. Thus, whatever the forces were that made EC the predominant marketplace location, they were not stronger than coordination effects. On 13 servers, some accident of history dictated that the marketplace would start up in Faydark or Freeport, and once it did, coordination effects ensured that it would stay there and not move to EC. And this happened even though, for example, evil characters had a much harder time participating in markets in Freeport than in East Commons. It happened even though the Faydark is located well to the east in Norrath's geography, placing it considerably off center in most travel routes. These two sites retained their status as the marketplace because, well, they were the marketplace. Coordination alone made their status robust. Thus, the evidence quite clearly supports the idea that coordination is itself a powerful macro-social force.¹⁰

The data can be parsed for more results that support the basic idea. For example, it seems as though EC was comparatively more predominant in the early stages of the game. Almost all of the servers founded in or before March 1999, the game's launch date, have markets in the EC zone (88%). The frequency of EC falls off considerably after that date (54%). Is it possible that changes to the game altered the relative advantages? Later evidence on this score is fairly persuasive: About a year after the time of the NES, the designers implemented a zone called the Bazaar and made it easily accessible from almost anywhere in the world. Moreover, in this zone money had no weight, banks were readily available, and good and evil characters were treated equally by the guards. Finally, in this zone players were allowed to automate their selling, setting up their characters as merchant robots whose wares could be searched through an eBay-style interface. No more hawking of goods and one-on-one negotiations, in other words. When the Bazaar was implemented, it was so obviously a better place to buy and sell that all of the old market zones immediately were emptied of all auctioning activity. At some level, market function must

dominate coordination effects. On the other hand, the opening of the Bazaar was not a secret—the designers acted very prominently as a third party, a voice sending signals to the entire population, telling them all “As of tomorrow, the market will be here.” Coordination theory predicts in that case that the market may well move if everyone’s mind can be changed at once. And in this case, it did, on every server. So perhaps the slight erosion in the dominance of EC early on has something to do with transactions costs, travel, or information.

On another dimension, the history of the game as a whole offers further confirmation of the role of coordination effects. As EQ grew, the designers created new servers. Sometimes they simply opened up new worlds for new characters only. Other times, they let players with advanced characters on older, crowded servers move those characters to the new server. The latter was called a split. Thus, on March 1, 2000, the server Fennin Ro was split into Fennin Ro and Druzzil Ro; Druzzil was a new server, just opened, and characters on Fennin could choose which server to play on. Meanwhile, because many servers were crowded, players from other servers would start new, Level-1 characters there too. A server split, in other words, created quite a mishmash of population on the new server, yet with a hefty chunk of advanced characters from the old server. If forces other than coordination were dominant—including this time something unique about the player base of the servers—then there ought to be some dependence in market location through these server splits. If not, if market location is driven mostly by random factors that then get reinforced by coordination effects, then servers that are born out of old servers should be unlikely to simply inherit the old server’s market location. Rather, each new server would be like a new experiment; the new population, whether transferring from the old server or starting anew, would replay the core experiment of market location and come to its own historical outcome, choosing its own location from among the foci of East Commons, Greater Faydark, and North Freeport.

Table 2 summarizes the seven server splits and their outcomes. There are seven splits, all from the very first servers founded in March 1999. All of those servers had East Commons as the marketplace. Yet after the split, only three of seven retained EC as the marketplace. Three of the new servers chose Greater Faydark, whereas one chose North Freeport. If some sort of direct force were dictating market location—specifically the force of remembered history, resident in the minds of players moving from an EC server to a new one—then every one of these servers should have had its marketplace in EC. Yet Faydark was just as likely as the Commons in this small sample. Thus, even though population is not controlled here, the evidence still suggests that coordination effects are a strong force in the placement of markets.

One last piece of evidence: The designers apparently intended for markets to occur, and it appears they had locations in mind. Every city had some sort of market area, spacious, safe, with nearby banking and nonplayer character merchants to facilitate trade. At times, these were even placed directly on trade routes, such as the tent merchants near the Port of Qeynos. However, with the exception of North

Table 2
Server Splits and Market Location

Date	Old Servers	Old Server Market Location	New Server	New Server Market Location
03/01/00	Fennin Ro	EC	Druzzil Ro	EC
04/28/00	Veeshan	EC	Saryrn	EC
	Bristlebane	EC		
07/19/00	Karana	EC	Terris-Thule	GF
	Bertoxulous	EC		
10/03/00	Mithaniel Marr	EC	Drinal	EC
	Tunare	EC		
11/21/00	Tarew Marr	EC	Xev	NF
	Povar	EC		
12/19/00	Cazic Thule	EC	Ayonaë Ro	GF
	Karana	EC		
	Solusek Ro	EC		
03/21/01	Rathe	EC	Zebuxoruk	GF
	E'ci	EC		

Source: Castronova (2001) and published information from EverQuest about server splits.

Note: EC = East Commons; GF = Greater Faydark; NF = North Freeport.

Freeport, it appears that actual human trade did not occur with any frequency in the places designed for it.¹¹ In the vast majority of cases, all of the constructed markets on all servers were as quiet as tombs almost all the time. This has two implications. First, merely designating something as a market is insufficient to force markets to locate there; although the voice of the third party can have an effect when choosing among equally prominent foci, it cannot make an entire population coordinate on a bad outcome. Things other than voice, such as transportation routes and local services, are more important. Secondly, this outcome makes EQ safe as an experimental testbed for a market location study. If the designers were attempting to induce some kind of market location outcome, they failed utterly. The location of the markets was an emergent property of the player base in its interaction with the world, unaffected by designer intent. It was as natural an experiment as one could devise.

In sum: The 40 servers of EverQuest observed in August 2001 constitute a reliable natural experiment on coordination effects in market location. The evidence indicates that the theory of coordination games does indeed operate on a large-scale level in human societies. Both the power and the fragility of coordination were apparent. If you were an evil character but lived on a Freeport server, the value of the market to you was lessened. Yet you still had to do your buying and selling in Freeport despite the hassle this apparently caused. The market was where it was; no one could easily change that; and its location had coercive power on all the players. At the same time, the fragility of coordination was evident as new split servers

seemed equally likely to move to a new market location as to retain the location of the parent server. In terms of both the breadth and depth of the theory, the markets of EverQuest offer rich evidence of coordination's existence and its force.

Evidence From Camelot

Mythic Entertainment's Dark Age of Camelot has received less scholarly attention than EverQuest, but for the purposes of this article they are equivalent: Both can be construed as large games with replication. As large games, they host macro-social effects no different from those encountered here in the outer world. As replicated games, they offer insights with the quality of controlled experiments. Like EQ, DAoC's server societies might tell us something about coordination effects. If so, such evidence lessens any doubts we might have that EQ is somehow special in this regard. That is important. The argument here is that the value of large games as research tools lies in their plumbing, in their structure. It should be possible to switch from a Tolkienesque world of elves and rings to an Arthurian world of paladins and grails and still obtain valuable pseudoexperimental insights.¹²

The evidence that Camelot brings to bear is best illustrated descriptively rather than through data; it is based on a number of points that were obvious to anyone spending time there.

The point of gameplay in DAoC was to engage players from other realms in PvP combat. Designers created three realms and then crafted three battleground zones, one belonging to each realm. Within each realm's zone were several strongholds and castles, the largest of which contained precious relics of the realm. The power of each realm's players was in part determined by how many of these relics their realm owned. Thus, players had incentive to enter the battleground zones of enemy realms, attack their castles, and try to take relics. At the same time of course, each realm had to defend its own relics.

The strategic situation from a general's eye view was an interactive game with no obvious outcome, much like rock-paper-scissors¹³: three zones, A-B-C, three players 1-2-3. If Player 1 sends his army to Zone A, and Players 2 and 3 do not, then Player 1 gets the relics in A (or keeps his relics from the others). If two or more armies met, there would be a battle, with the winner getting relics. In any round of play, there are multiple possibilities: no conflict, 1-on-1 conflict, 2-on-1 conflict, or 1-on-1-on-1 conflict. On the surface, this game looks to provide the same kind of fluidity of outcomes that rock-paper-scissors does except that, as with poker, knowing the tells of the other generals could be extremely helpful. Assuming a reasonably open promotion system within the player bases of the three realms, one could guess that the three leaders would be pretty good generals and not have obvious tells. Thus, one would expect balance among the realms and changeable, exciting combat going on each night in the three battleground zones.

However, as Camelot aged it emerged that most players spent most of their time in only one of the battleground zones, Emain Macha. It was not entirely uncommon to encounter organized combat in the other zones, but raids on relic castles were fairly rare. And the fighting in Emain was rarely organized. Rather, the forces of Albion and Midgard—technically, invaders—would mob together at their respective invasion points, which were quite close, and the forces of Hibernia—technically, defenders—would amble the long distance from their homeland out to this frontier. Thus mobbed in proximity, the players would engage in more or less endless mob-like fighting—a thrust here, a retreat there, a panic, a few foes vanquished; run back to safety and try to repeat. Although this was “victory,” no relics changed hands. The Emain *zerg*, as it was called, was considered by many players to be a letdown in terms of gameplay, but nonetheless it was the primary PvP pattern produced by Camelot’s design.

A number of factors explain the deviation of player behavior from designer intent, and coordination game theory has a dominant role. Let’s begin with some critical details. First, Emain had wide green fields and gentle slopes in contrast to the other two battlegrounds, which were mountainous and densely forested. (One player described Emain as “a big golf course.”) Second, the rewards of relic taking were shared among the entire population of a realm, whereas rewards from killing foes in individual combat, called “realm points,” accrued to the individual player. Third, the game offered no explicit or implicit authority structure among the players other than a system of self-organizing and voluntary player clubs (“guilds”). Fourth, there were no supply lines, no sense of the front and rear or a fighting force, or any way for one body of troops to impede the movements of an enemy force—one player could simply run through another as though that player did not exist. Combined, these features turned the combat game of Camelot into a coordination game, with extraordinarily powerful incentives built in to play that game one and only one way.

To begin at the end: With no authority, no supply, and no sense of military space, there was no reason in Camelot for armies to even exist. No one could actually give orders and expect them to be followed; no tools existed whereby this kind of authority could emerge. On the contrary, the tools effectively prevented the emergence of authority by making sure that anyone could leave any group, guild, or military formation at any time and face no worse consequence than the annoyance of their erstwhile comrades in arms. Meanwhile, the absence of contact detection, supply, and so on made sure that no one had a self-interested incentive to stick together or hold a particular formation or piece of ground. This almost forced all fighting formations to be loosely affiliated mobs of players. This then was reinforced by the incentive scheme. Public choice economics (Buchanan & Tullock, 1962) long ago clarified the simple fact that rewards shared by a collectivity are dramatically less motivating than rewards one obtains for oneself. If individual rewards are obtained for combat pure and simple, whereas relic raids give rewards to the entire realm, it is obvious that most players will spend their time seeking out combat for combat’s sake alone and

leave relic taking to special circumstance. Thus, we have loosely affiliated mobs of players, all of whom seek to accomplish nothing more than combat with one another. And lest this seem to be a harsh critique of the designers, it was obvious that the players found this kind of fighting extremely fun. It is not fun to march about for hours while the generals maneuver, only to be ordered at the moment of battle to stay back and guard the baggage. Not fun. Much more fun to engage in endless random melee that has no real purpose or consequence for anyone. Because it was flexible, allowing both relic raiding and zergwar, Camelot's design should be judged a success. Nonetheless, it seems that the predominance of zerging was a surprise to everyone.

Given these incentives to zerging however, coordination then plays a huge role in making it happen as players want. Here we have mobs of players who want nothing more than to attack one another, a lot, and so rack up thousands and thousands of realm points. To do that, the players of the three realms do best to identify one spot for fighting and go meet one another there, always. The worst disaster would be to have confusion about the location of the battlefield. What if the Britains went to the Viking home of Midgard, the Vikings to the Celt home of Hibernia, and the Celts to the Britain home of Albion? No zerg and no points. It behooves all players, as with the market location example, to settle on one spot. And in this respect, the lush, rolling fields of Emain were evidently preferable. Enemies could be easily spotted and closed to. Anyone looking for a mob of fellows could easily find them. Meanwhile, rogues and other invisible characters who were generally considered impossible to beat if they got close while still hidden had fewer places to hide. In fact, the only argument against Emain would have been made by the Hibernians, who had to travel much farther to get to the fighting areas than the other two realms. Yet the advantages of the space and its easy affordance for zerg combat overcame the natural resistance of Hibernia. Thus, on each of Camelot's 20+ servers, most combat occurred in Emain Macha because of the combination of player incentives and a corresponding powerful return to broad social coordination.¹⁴

Here again we have evidence of the power of coordination effects, and again the evidence is richer because of the quasi-experimental nature of the large game in which these effects were observed. If coordination effects were not powerful, then the Hibernians should have been able to force changes from time to time in the location of the zerging warfare on at least some servers. They were not. Rather, Hibernia was everywhere forced to accept the benefits of coordination with Albion and Midgard despite the cost of doing so in Emain Macha. And later developments suggest that the coordination effects were also fragile. In a later change in the geography of the zones, the Hibernian battleground was no longer made into a flatter, opener version of the others, and the concentration of warfare there ceased.

Thus, here we have a second example of a large game serving as a natural experiment on the subject of large-scale social coordination. If a designer establishes 20+ versions of the same world and imposes on each version a similar set of incentives to mob warfare and then creates one zone where mob warfare is particularly easy to

accomplish, he or she has built an experiment where 20+ human societies are playing a coordination game and one of the game's focal points has a higher payoff than the others. Theory predicts that each society should gravitate to the high payoff focal point, and that is in fact what has happened. This is a second experimental confirmation of the existence of coordination effects at the macro-social level made possible only by the uniquely (though naturally, through no intent of the designers) controlled conditions presented by a large game.

Conclusion

This article set out to make two points:

1. Coordination effects really do happen in large human societies.
2. The best way to find evidence of them is in large games.

Support for these points relied on a common intuition about the nature of human society: It runs on some kind of underlying structure, a structure that operates in roughly the same way regardless of where the society is located in space or time. The structure of social interaction does not predict all things social, but it does predict some of them. Thus, supply and demand and the invisible hand have something to say about the prices of goods no matter where those goods are or what other things are used to price them. I've argued that coordination effects are similarly a common feature of human society, that under certain circumstances we find ourselves in a huge game where there are on the one hand powerful incentives to conform to the expectations of others, while on the other hand, it is obvious that merely by changing our expectations collectively we could cause our behavior to conform to some other norm. And if coordination and conformity operate at some abstract level, they should be observable in all kinds of societies, including the societies contained in large games. Large games are large enough, conceptually and technologically as well as in terms of raw population and hours of use, to be considered genuine human societies in their own right. Unlike most human societies however, the societies of large games can be grown, as it were, in carefully controlled circumstances. And one circumstance that has been unintentionally built by the designers of several large games is replication: A single large game with a single structure is built over and over again in exactly the same way. As such, the set of replicated worlds represents multiple trials of an experiment. And in these trials, we see things that no amount of historical, theoretical, or deep-knowledge research can show us: We can see the probability distribution of a human society, viewing its features not as from a single outcome but from the whole range of things that might have been but were not. The opportunity to see large-scale social outcomes from a truly probabilistic, experimental perspective as in a petri dish has never before existed in the long history of social thinking. It is an advance of incredible power and value.

I argued that large games, though valuable as petri dishes, are also expensive and thus represent a social science version of large-scale physics tools like supercolliders. To demonstrate the value, I explored natural experiments provided by two social science supercolliders that already exist (though unwitting in their role as tools of research): EverQuest and Dark Age of Camelot. These two behemoths of gaming, each with more than 200,000 active subscribers, have in pursuit of commercial objectives crafted controlled circumstances for bashing atoms of social science theory into one another, crudely, and without reference to any research hypothesis whatsoever. Yet we can observe the trace patterns left by these collisions, and they tell a reliable story: Coordination effects really do occur at the level of all society. When the incentives are right, a whole society of real people act as coordination theory predicts. This result is based on direct observation of dozens of comparable, parallel societies. It is important to contrast this directly with the methods currently available to social scientists. To make this absolutely clear,

- The results are not based on a study of coordination among 12 to 16 undergraduate students in a lab playing abstract games for a few dollars.
- The results are not based on an analysis of coordination by some groups in France as compared to some groups in Germany or in France in an earlier age.
- The results are not based on a pure theory of what coordination might be without any reference at all to actual human behavior.
- The results are not based on the researcher's impression after having spent 12 months living with a small subset of one of the populations.
- The results ARE based on the direct observation of parallel, replicated societies in which coordination behaviors were possible but not induced.

It should be apparent from the tone that I think the latter mode of study is at least as reliable, and quite probably more so, than those that precede it. Whereas the evidence from experimental large games does suggest that coordination effects exist at the macro-social level, I would argue that there is a bigger finding here. The big news is not that coordination effects exist but that large games can find them. That being the case, a major realignment of social science research methods would seem to be in order.

Notes

1. There is a vast literature of purely theoretical treatments of coordination games. Interested readers should start with basic textbooks, such as Fudenberg and Tirole (1991) or Dixit and Skeath (1999).
2. To my knowledge, there are no comparative studies that invoke coordination games as a theory for explaining differences in outcomes. The economic historian Avner Greif (1993) however used coordination theory to study trade patterns in the Mediterranean.
3. There is a large literature on experimental games. Again, the interested reader would be directed to the leading texts in game theory (see Note 1).
4. I know of no ethnographic studies that invoke coordination game theory to explain observed behavior.

5. Recent research in psychology disputes assertions that the things in Norrath are “not real.” A significant body of research indicates that the mind treats virtual things largely as if they are real (Reeves & Nass, 1996). The labels of *real world* and *fantasy world* are apparently mere conveniences for most game players—they just identify where something happened, not whether it has any existence or meaning.

6. I am not saying that a theory has to make the same predictions everywhere or that all societies are the same. On the contrary—a good theory allows for and explains social diversity, just as the theory of supply and demand explains why the price of rice is higher in Moscow than in Bangkok. Moscow and Bangkok are very different places, but through the theory we can see why they are different and what effect their differentness has on things we can observe, like prices. My point is only that a good theory has to work with nuts and bolts that have broad relevance.

7. For those not familiar with the structure of massively multiplayer online role-playing game (MMORPGs): Technology constrains the maximum number of players that can share a given world at the same time. At the time of this writing, the constraint is about 3,000 or 4,000 players. Game companies know that if 10,000 users have an account with access to a space, only about 3,000 or so will be online at any one time. Therefore, one scales a game like EverQuest (EQ) on a server by server basis: The first 10,000 subscriptions go on Server 1. If more than 10,000 people sign up, Server 2 is opened. If more than 20,000 sign up, Server 3 is opened. And so on. Thus, EQ’s 400,000 players were hosted on 40 servers, 10,000 players each. The critical point is that each of these servers is providing exactly the same game to all the players. The terrain is the same, the enemies are the same, the quests are the same, the transportation is the same. At most, there might be slight variations in the rules (as when players are allowed to use their combat abilities against one another: player vs. player, or PvP, which is not the norm). Thus, each server is essentially a replica of Server 1: another in a long line of large-scale, societal Petri dishes.

8. To perhaps better express how poignant such questions have been, consider this: What if we reran German history 40 times? Would Hitler have advanced to power in every incarnation of the Weimar Republic? In most of them? In only a few? Was Hitlerism a rare and tragic mishap for the world? Was it avoidable? Unlikely? Or predetermined, fated, inevitable? It is hard to overstate how valuable that information would be from the perspective of emotions no less than policy to hundreds of millions of people living today. If we had 40 Weimar Republics, we could make a judgment about those questions.

9. The reasoning behind each of the three marketplaces is as follows. East Commons (EC): The market there took place in the EC Tunnel, a tunnel that allowed players to move from the East Commons zone to the North Ro zone. The tunnel offered many advantages as a marketplace. First, here good and evil characters could freely interact, whereas in most cities the one or the other would be chased and killed by guards. Second, East Commons itself was a crossroads for both types. Good characters crossed through here when passing from the human cities of Antonica, in the west, to the elven lands of Faydwer in the east. Evil characters came through East Commons when traveling between the ogre and troll lands of South Antonica and the Dark Elf empire of Neriak in North Antonica. Third, the EC Tunnel was spacious and had clear landmarks. On the walls were a series of torches, and merchants would announce location by saying “see me at the second torch.” Finally, the EC tunnel had no monsters at all and was completely safe. Greater Faydark (GF): Although located well to the east, GF was home to two of the most popular races, the wood elves and high elves. Moreover, two other races, gnomes and dwarves, would cross through GF to get to adventuring areas. GF is a dark place with fairly low-level monsters and lots of open spaces, meaning that good and evil characters could freely mix and mingle. It was home to both a druid ring and a wizard spire, which at that time meant, in effect, airport: At those sites, druids and wizards, respectively, could transport players to and from far-flung locations. Finally, unlike EC, GF had convenient banking services. North Freeport (NF): NF was a zone inside the city of Freeport that was adjacent to East Commons. As such, it benefited from a similar crossroads effect and being a city zone, there were no monsters at all. Like GF, NF had a convenient bank. Unlike EC and GF however, NF was not congenial to evil characters; evil characters had to implement workarounds to do their marketing in that zone. On the other hand, of the three candidate zones, NF is the only one that looks like a marketplace, with stalls and nonplayer character merchants scattered about.

10. I would venture to claim that this is the first time in history that a distinct macro-social phenomenon has actually been verified experimentally.

11. The elf city of Kelethin was in the Greater Faydark zone, so one could technically say that the GF market was “in” Kelethin. But practically speaking, it was not. Kelethin was a city in the trees, and no one went up there to do their trading.

12. There are actually a number of concrete differences between these large games at the level of plumbing as well. Three worth mentioning: EQ allows PvP only as an exception, whereas in Camelot it is the rule. In EQ, all players interact with one another, whereas in Camelot the player base is split into three factions that have no technical means for communicating. Finally, in Camelot the level of market activity was lower than in EQ, mostly because fewer goods were sensibly tradable.

13. Formally speaking, the assertion is that like rock-paper-scissors, the game does not have a pure strategy equilibrium.

14. Here I do not have survey data to back up the assertion that Emain was a zerg combat locale on all of Camelot’s servers. Yet I think it would be undisputed by any open-minded visitor to Camelot that Emain was known throughout the lands in a nearly universal opinion to be the location of mindless zerg fighting every night on every server.

References

- Buchanan, J. M., & Tullock, G. (1962). *The calculus of consent: Logical foundations of constitutional democracy*. Ann Arbor: University of Michigan.
- Castronova, E. (2001). *Virtual worlds: A first-hand account of market and society on the cyberian frontier* (CESifo Working Paper 618).
- Dixit, A., & Skeath, S. (1999). *Games of strategy*. New York: Norton.
- Fudenberg, D., & Tirole, J. (1991). *Game theory*. Cambridge, MA: MIT Press.
- Greif, A. (1993). Contract enforceability and economic institutions in early trade: The Maghribi Traders’ Coalition. *American Economic Review*, 83, 525-548.
- Malaby, T. (2005). *Ethical code: Modernity and the game concept in computer game design*. Working paper.
- Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, televisions and new media like real people and places*. Cambridge, UK: Cambridge University Press.
- Schelling, T. C. (1960). *The strategy of conflict*. Cambridge, MA: Harvard University Press.

Edward Castronova (PhD, economics, Wisconsin, 1991) is an associate professor and the director of graduate studies in the Department of Telecommunications at Indiana University, Bloomington. He is an expert on the economies of large-scale online games and has numerous publications on that topic, including a book titled *Synthetic Worlds: The Business and Culture of Online Games*.