INTRODUCTION:
THE CHANGING MEANING OF PLAY

Why Are You Paging Through a Book about Games?

According to the jacket material, this book claims to take a serious look at online computer games. You are actually reading it, at least the first few lines of the introduction, and if you are a serious, hardheaded person like me, you must be feeling a bit strange. I feel a bit strange writing it. Three years ago I was an ordinary economist pursuing generally ordinary economics research. Now, however, I am pushing deeper and deeper into a realm of experience that's growing faster than I can examine it, a fantastic cosmos of dragons and rayguns and beautifully crafted human bodies. It is also a universe that hosts massive flows of real human intercourse—information, commerce, war, politics, society, and culture. I am speaking, of course, of the phenomenon known as "massively multiplayer online role-playing games" (MMORPGs), places where thousands of users interact with one another in the guise of video game characters, on a persistent basis: many hours a day, every day, all year round. As such, these places are like real cities and fairy-tale cities at the same time, and some of the numbers they are producing might surprise you:

- Users drive around in these worlds using a video game character in much the same the way we use a car to drive around the Earth. Some characters are better than others: faster, better looking. They can be bought and sold, most often on eBay. As I write this, a Jedi-type character from a fantasy world based on *Star Wars* costs over $2,000.
- Typical users spend 20–30 hours per week inside the fantasy. Power users spend every available moment. Some 20 percent of users in a recent survey (see

chapter 2) claimed that their fantasy world was their "real" place of residence; the Earth was just a place you go to get food and sleep (see chapter 2).

- This used to be a niche phenomenon. But synthetic worlds are appearing at the rate of Moore's Law (i.e., doubling every two years), and the current number of users is 10 million people, at a minimum (see chapter 2).
- Each synthetic world has a play-money currency inside to facilitate player-to-player transactions. These currencies have begun to trade against the dollar in eBay's Category 1654, Internet Games. Many of them now trade at rates higher than those of real Earth currencies, including the yen and the Korean won.
- The commerce flow generated by people buying and selling money and other virtual items (that is, magic wands, spaceships, armor) amounts to at least $30 million annually in the United States, and $100 million globally (see chapters 4 and 6).
- In Asia, people who have lost virtual items because of game-server insecurities and hacks have called the police and filed lawsuits. The police have made arrests; courts have heard cases; and plaintiffs have won (see chapters 6, 7, and 11).

It is hard to look at these developments without concluding that something quite bizarre must be going on, and perhaps something just as important as the subjects I used to work on as an ordinary economist. Thus while I have little doubt that I am just as ordinary as I was before, the subjects I write about have become quite extraordinary. Indeed, *this* subject, video games, is so extraordinary that it has attracted the attention of one such as yourself—a person who does not normally think of video games as anything more than child's play. The aim of the book is to change that view. If you had seen what I have seen in the last few years, you would sense, as I do, that the line between games and real life has become blurred. Together we might begin to understand how much this blurring will change the nature of daily life for our children and grandchildren.

The thesis of the book is that the synthetic worlds now emerging from the computer game industry, these playgrounds of the imagination, are becoming an important host of ordinary human affairs. There is much more than gaming going on there: conflict, governance, trade, love. The number of people who could be said to "live" out there in cyberspace is already numbering in the millions; it is growing; and we are already beginning to see subtle and not-so-subtle effects of this behavior at the societal level in real Earth countries. Even if you haven't paid much attention to multiplayer video game worlds up to now, soon enough, I think, you will. We all will.

Perhaps the easiest way to convince yourself that this hypothesis is, or is not, true would be to go visit one of these places yourself. Flip through the book; certain names will come up more than a few times: *Ultima Online*, *EverQuest*, *Lineage*, *Second Life*, *Dark Age of Camelot*, *Star Wars Galaxies*. Or visit your local game store; as this book goes to press, several hundred thousand people are exploring the brand-new *World of Warcraft*. You could find the software, pay for subscriptions, and head out into these places to see what may be seen. I predict, however, that you might find this means of discovery awfully expensive in terms of time. You would be trying to study a completely different culture, one for which we have no prior literature and no guidebooks. The natural first reaction would be bewilderment, of course. It was for me. Only after hundreds of hours of immersion did I begin to have any success understanding what is happening there and what it might mean.

Part 1 of this book attempts to shorten the process of discovery by taking the reader on a guided tour of synthetic worlds. I've found that for many audiences, mere exposure to this phenomenon as it is today is sufficient to render obvious some of the deeper consequences, for those aspects of human life and thought that the listener knows most about. Anthropologists see new cultures, entrepreneurs see new markets, lawyers see new precedents, and social and political experts see new pressures and looming crises. Part 1 attempts to deliver a basic understanding of what this technology is, how it works, how people use it, and the kinds of social institutions that arise because of it. At that point, the long-run implications—and as a corollary the merit of video games and interactive media as objects of serious study and reflection—may already be apparent. We do not have to look very far into the future to see changes looming.

As a quick test, imagine someone told you that there was a technology that could reasonably be referred to as "practical virtual reality." This technology would allow just about anyone, at a modest cost, to spend as much time as they wished in some kind of alternate reality space that was built and stored on a computer. We are not talking about a Holodeck here; this place isn't "real" by any means.[1] However, it does feel real enough to the users that they can fairly easily immerse themselves in it, for hours on end, month after month, year after year, in a sort of parallel existence. Were such a practical virtual reality possible, we can have no doubt that profit-seeking enterprises would figure out a way to make it work. They would build around certain themes and construct entertaining activities, to draw people in. How many people might be drawn into such places for hours and hours on end? Imagine there were a large number of just-good-enough fantasy worlds for people to go live in, worlds with all kinds of themes, from knights in armor to athletes to space travelers to mobsters to almost-credible lovers. How many people?

The fact is, this "just-good-enough" virtual reality technology exists today. By the time you read this, we will have already moved beyond it, to "almost-

seductive." And if you didn't answer "just a few" to the question that closed the preceding paragraph, you might want to take a break here, and then at the end of part 1, and ask yourself how things might soon be changing in your area of interest.

As I argue in part 2, it is not hard to imagine that there will be major effects in many areas. The overarching idea in the second half of the book is that the emergence of these practical virtual reality spaces will have significant consequences primarily because events inside and outside them cannot be isolated from one another. It is not too shocking to imagine that the real world can affect the virtual world; when it rains on a football game, the game is changed. But we are now learning that games may become so important to some people, at some times, that events inside games have effects outside of them. Should more people become involved in practical virtual reality spaces, these external effects will become quite serious on a macro level. While one could make a case for these external effects in a number of areas, especially culture, sociality, relationships, and individual emotion, I will focus more narrowly on my own areas of interest in the social sciences: economics, politics, and security. In my opinion, major changes in these areas would be quite serious indeed, well worth contingency planning. Part 2 of the book is primarily concerned with assessing the potential for such changes.

Then in the final part of the book I draw broader conclusions for public policy. Again, imagine you had mastered a technology of practical virtual reality. What would you do with it? It quite clearly could be good or bad for humanity, depending on its usage. Given that the technology does exist, and is being driven onward by a lucrative, savvy industry, we should be aware that while we may have quite a few years to think about all this, it behooves us to start thinking right now, before the time to make important decisions arrives.

What Is a *Synthetic World*?

So far I have been referring to the technology in question as a "practical virtual reality" tool, a way to make decently immersive virtual reality spaces practically available to just about anyone on demand. For the most part, I will refer to these places as *synthetic worlds*: crafted places inside computers that are designed to accommodate large numbers of people. The specific incarnations that I will talk about in this book are places created by video game designers. Chances are, if you are under 35 years old you know exactly what that means, because you have been playing in synthetic worlds since you were a kid and you know that they have moved online. If you're not part of the video game generation, you might have

trouble seeing the connection. Your kids or grandkids play video games on their TVs, on desktop computers, on little handheld devices—what does that have to do, you might ask, with "virtual reality"?

Such a question makes good sense because games and handheld video devices are not part of our standard image of what virtual reality is all about. For most people, I suspect, the first thing that comes to mind when one thinks of "virtual reality" is a laboratory filled with expensive gear. You know, the bulky vision-goggles that go over your head, the web of wires and straps for your arms and legs, the funny half-chair, half-bike apparatus that you wiggle around on, the six-sided surround-sound rooms at Disneyland that make you feel like you're standing in the Amazon jungle or on the Moon. Jennifer Lopez in *The Cell* (2000): bodies suspended on wires, wrapped in a second skin that looks somewhat like ribbed beef jerky. *That's* virtual reality, right? Well, no. At least, not entirely, not any more. At one time, virtual reality was indeed a matter of basic lab research and ingenious sensory-input devices, a scientific research program that made headlines in the early 1990s. But the virtual reality I am talking about has emerged independently of that program; it grew out of the game industry, without any influence from the scientists. Game developers had been exposed to the same basic ideas of virtual reality that everyone else had—Gibson's *Neuromancer* (1984), Vernor Vinge's *True Names* (1981), and so on—but they took them in a completely different direction. The difference was this: the science program focused on sensory-input hardware, while the gamers focused on mentally and emotionally engaging software. As you can imagine, a person can become "immersed" either way: either the sensory inputs are so good that you actually think the crafted environment you're in is genuine; or, you become so involved mentally and emotionally in the synthetic world that you stop paying attention to the fact that it is only synthetic. It turns out that the way humans are made, the software-based approach seems to have had much more success. It certainly is more popular, and also cheaper for users and developers. And so, as we head into the twenty-first century, the dominant paradigm for virtual reality is not hardware but software, and that means that any device, even a crude one, that can engage a person in the happenings it portrays, is a little virtual reality tool. When children play at their little handhelds and when executives fiddle around with the games on their smartphones, there's immersion going on, a virtual reality brought about by games rather than devices.

Now, the contrast between the scientific and the gaming approach to virtual reality is important to the intellectual placement of a book like this, but it's also something of a tangent, mostly because the interesting thing about these two fields is how little they have to do with one another. There's a man-bites-dog story here: virtual reality lost much of its status as an exciting technology over the past

decade, and while it is now re-emerging with considerable force, what's strange is that it is not emerging from the ruins of the old paradigm. Rather, it is growing up in a completely different place. The story of how these development tracks emerged independently of one another is worth telling, at least for those interested in virtual reality in general and its intellectual history. However, the target audience for this book is not the virtual reality specialist but rather professionals in business, government, and education, parents, journalists, and academics, basically folks in other walks of life for whom this new and practical virtual reality represents genuine opportunities as well as threats. The members of the Army Air Corps in 1918 were not so interested in the fact that the bicycle makers who built the first airplane modeled it not on the vehicles they had been building before, but on something completely different, the automobile. That's fascinating history, but it's not necessary to an assessment of the technology's impact. No, people who are going to be affected by new machines just want to know what they are all about and what they can do.

So I won't be spending much time on the traditional virtual reality paradigm, other than to describe, as a lengthy aside (see the appendix at the end of the book), why its vision did not lead to the massively populated, yet mechanically crude, virtual reality spaces we have today. Readers interested in the contrasts between paradigms could take a pause and have a look at the appendix. Those not interested in that history could continue here without significant cost.

For those who did not grow up with video games and are not all that familiar with previous efforts to construct virtual reality, the basic idea behind the video game as a practical virtual reality tool is this: If the game is online, a user can log into it from any computer on the Earth. The screen turns into a window through which an alternative Earth, a synthetic world, can be seen. This other place (another planet, a historical domain, or any other plane of existence) can have mountains, stars, and fire in it; it can have gravity, or no gravity, or reverse gravity; it might have trees and grass, but also chickens and dragons, or chicken-headed dragons or dragon-headed chickens; it might have houses and taverns and castles, or spaceships, or tiki bars; and it might have people. Some of the people you would see might be software-controlled, but others would be controlled by real humans, such as yourself. In fact, there might be a mirror there, and if you press the right buttons and maneuver your viewscreen in the right way, you would see yourself, present, in that place. The window by which your computer is depicting the world is, in fact, the surface of somebody's eye, and that somebody is *you*.

More accurately, you have been given a synthetic body in the synthetic world, and your computer is rendering the world as it would be perceived by the ocular

sensory device that your synthetic body possesses. If you see someone else in the world, and she is pointing her visual sensing device at you, well, the two of you are looking at one another through your computer screens. She may be in Hong Kong, and you may be in New York, but you are still occupying the same segment of cyberspace and you have just made eye contact. Of course, her appearance there may not match the looks of her body on Earth, but neither does yours. She may even be a he, not a she, or something else again. But she/he/it is a person, like you, and you can have a relationship with her that is just like any other relationship you might have with another person. The only difference is that this relationship is being mediated by a body that is one step removed from the Earth body, and therefore occurs in a place that is one step removed from the Earth.

I said above that synthetic worlds are becoming important because events inside them can have effects outside them. This flow of influence from inside to outside is generated by a very simple core mechanism that is easy to see in the vignette above. As soon as it goes online and begins to receive visitors, a synthetic world begins to host ordinary human affairs. However fantastical the place may be—whether visitors are represented as mobsters, dragons, or crumb cakes—it still and always is playing host to ordinary human beings, with their ordinary ways of interacting with one another. The physical environment is entirely crafted and can be anything we want it to be, but the human social environment that emerges within that physical environment is no different from any other human social environment. And because no one can permanently separate events in one sphere of their life from all the other spheres, that part of human life taking place in synthetic worlds will have effects everywhere. At the same time, the things that happen there will not be run-of-the-mill things. We will no longer be in Kansas, and many of the rules will be different. Thus not only will there be spillover effects; the effects, such as they are, will seem weird.

In short, synthetic worlds put ordinary humanity in a very strange place, producing forces that deserve hardheaded attention, in my view. All things that matter to ordinary people—their loves, their crusades, their morals, and their material assets—may now have a home in a place other than Earth. That place operates under different rules. As Lawrence Lessig (1999) describes it, the unusual thing about cyberspace is that we can be both here and there at the same time, and the place that is "there" can be constructed, essentially, however we might like. Thus, all of our interests are the same as they ever were, but the environment in which we pursue them has become untethered from the Earth environment with which we have become so comfortable.

Deep questions arise, only a few of which we will have space for. But it is easy to find a simple example of the way that the strange and ordinary features of

synthetic worlds can conspire to generate phenomena worth thinking about. Consider the price of diamonds. On Earth, these items tend to be quite expensive. They are beautiful, and if crafted well, they enhance the beauty of those who wear them. Their beauty contributes to their price, of course, but so does their scarcity. Now, what if the Earth could be induced to produce as many diamonds as anyone would ever want? Such a thing is impossible here, but not in cyberspace. The coding authority who owns and controls a synthetic world could pave the streets with diamonds if it desired. The coding authority could also order the world to make gems that are more beautiful than diamonds. Or, the coding authority could make diamonds much rarer than they are on Earth, and then also purge the world of rubies, emeralds, sapphires, and every other type of gem. All of these coding decisions would affect the price of diamonds and the happiness of the people wearing them. The role of diamonds in society would be affected: Would the lover still give the beloved a glittering gem if it becomes common, or impossibly rare? How uncommon must a diamond be to serve as a love token? Or as a power symbol? Such questions are moot on Earth, because our planet is endowed with a certain availability of diamonds based on their presence in the ground and our understanding of how to get them out of the ground. In synthetic worlds, things are different. The availability of diamonds is not an endowment but a choice. Thus while the mental objects in play there (beauty, price, love, profit, scarcity, reputation, power) are nothing new, the rulebook under which they are all contested is a new thing indeed. Moreover, a love token is a love token is a love token: If the coding authority messes around with its virtual diamonds, it will affect some aspects of the love relationships between living, breathing humans. The heartaches and joys that result will not be virtual at all, nor will the behavioral reactions on Earth, which may range from new sleeping patterns to substantial purchases of alcoholic beverages. These will be genuine feelings and actual phenomena, as real as anything else under the sun.

Thus, it is the core features of synthetic worlds—the fact that they are radically manufacturable places that can be shared by many people at once—that generate their nontrivially unique inside-to-outside patterns of influence. Place a group of people in a strange place, and they will follow their usual tendencies in pursuit of their usual objectives. The outcomes will be both strange and familiar, and will radiate outward.[2]

Moreover, this combination of the bland, the fantastic, and the consequential points to the simplest answer to the question of what synthetic worlds really are: a frontier. Indeed, viewing synthetic worlds as a locus of migration is useful in coming to practical terms with their effects, and there certainly is some evidence of motion from "here" to "there." Statistics reported in this book will suggest that

many people are diving into the new worlds right now, with enthusiasm. Evidently, they find the physical environments crafted by computer game designers much more attractive than Earth. Accordingly, these travelers or colonists have come to maintain a large fraction of their social, economic, and political lives there.

These migratory patterns accord with the predictions of cyber-theorists such as William Gibson (*Neuromancer*, 1984) and Neal Stephenson (*Snow Crash*, 1992), who saw no long-run limit on the amount of time people would want to spend in virtual reality environments, were they to become practical. Now indeed they seem to have become practical enough, in a rudimentary way, to serve as way stations between the late twentieth century and the future as these authors envisioned it. Furthermore: It has already been several years since perceptive people like John Walker (1988), Michael Vlahos (1998), Ray Kurzweil (1999), and Hans Moravec (1999)—none of them novelists—began to argue that that future was nearer than we might think, that humans and their machines were going to become more intimately involved with one another, and fairly quickly at that. The natural location for getting together is in cyberspace, of course, through simulated bodies in simulated spaces. Why would you type into your computer if you could talk instead? And why talk to a gray screen when you could instead have a conversation over a virtual lunch with an attractive humanoid-looking being? These kinds of interactions, which could occur most easily in a crafted, synthetic world environment, seem quite plausible to anyone thinking about the growing power of computer technology. From this point of view, synthetic worlds are simply intermediate environments: the first settlements in the vast, uncharted territory that lies between humans and their machines.

Names for a New World

Today, of course, most synthetic worlds are considered games, and the term used to describe them—prepare your tongue, this won't be easy—is *MMORPG*. This jawbreaker (I pronounce it "mor-peg") emerged from a game industry practice that refers to games like *Dungeons and Dragons* as "RPGs," for "role-playing games." Specify an RPG as an online game, and it becomes ORPG. Such a game with multiple players is a MORPG, "M" standing for "multiplayer." Around 1996, the industry acquired the technical capacity to expand the number of players from a then-ordinary number (8–16) to what was considered a very large number (3,000–4,000), and the term "massively multiplayer" was coined. MMORPG, for "massively multiplayer online role-playing game," has become the standard term

of reference for all synthetic worlds—a development only the Pentagon (and some university administrators) would admire.

There are other words to use, even some that can be pronounced. The ancestors of MMORPGs were text-based multiuser domains (MUDs), commonly referred to as *virtual worlds*. The usage has now been applied to a wide variety of Internet spaces, but because of the unfortunate history of the "virtual reality" scientific research paradigm (see the appendix), the "virtual" tag has lost its meaning in many respects. Vlahos refers to the terminus of our migration as the *infosphere*, but he is talking about networked communication in general. William Gibson coined *cyberspace*, a term that has become too general, unfortunately. Neal Stephenson used the term *metaverse*, but that concept does not reflect the rendered, role-playing aesthetic that we now see is an important aspect of synthetic worlds, and by its derivation from *universe* it overlooks the fact that there are going to be many thousands of synthetic worlds to choose from, not just one.

For a more targeted term, consider *proskenion*, from ancient Greek theater. Brenda Laurel (1993) and Janet Murray (1997) propose that much of what happens in computers and online is a form of interactive storytelling or theater. Theater in Greece began as religious ritual and took place in a circular area at the foot of a hill. When ritual became drama, a small tent or hut called the *skene* was built at the back of the circular area, to facilitate costume changes and exits and entrances. Eventually, visual imagery (scenery) was painted on the *skene*, and then a raised platform was constructed on its front. This structure, the *proskenion*, is arguably the first physical space constructed by humans explicitly to serve as some *other* place, indeed, a place that exists only in our imaginations. In contemporary theater, *proscenium* is a technical term referring to the area on the stage that is outside the curtain but in front of the audience. This is apt as well. It suggests that "proskenion" serves as a term for any rendered, imaginary place, and it also reminds us, metaphorically, that the stages on which online dramas play out involve us both as actors and audience—the action occurs in a place that is not quite the stage, but not quite the seats either.[3] And once the audience becomes the players, of course, the play is no longer a play; it's ordinary life, even though it happens on a stage.

Indeed, following the theatrical metaphor, another term that might make sense is *hyperstage*. When programming allowed text to deliver radically new levels of information, text became hypertext. Similarly, MMORPGs take the concept of an ordinary stage and extend it radically. These places are theatrical locales in the sense that they are not of the Earth, and in that they allow people to take on many roles (including their true self if desired). But they are much more than that. They

allow such a huge number of players, and such an unscripted plot, that the line between acting and mere living blurs and, in many cases, vanishes outright. Moreover, we can teleport from one stage to the next with the click of a button, just as hypertext takes us from page to page instantly. Foreshadowed in Shakespeare's *As You Like It*, a hyperstage is a stage of the whole, a place in which computing technology has erased the distinction between actor and audience, here and there, scenery and landscape, role and self.[4]

While there might be a number of useful new terms, I will stick primarily with the term *synthetic world*: an expansive, world-like, large-group environment made by humans, for humans, and which is maintained, recorded, and rendered by a computer.

The Consequences of New Frontiers

However we refer to these territories—virtual world, MMORPG, cyberspace, metaverse, proskenion, hyperstage, or synthetic world—the most general causes and effects of any migration into them may not be hard to predict. Human migration is a well-known and fairly well-studied phenomenon (Borjas 2001). A simplified economic story would say that those doing relatively less well in one place face the risks of change and head off into a new place. They stake claims there but retain ties with their former neighbors. If they do well, they stay; if they don't, they go back. Economists argue that as people sort themselves into different places based on their skills and preferences, everyone is better off. Those who do well by moving, move; those who do well by staying, stay; and everyone eventually finds the best possible place to be.

While this is a happy story in the long run, nonetheless, it is also a story of great change and short-run stress. Social and cultural attitudes, technology, and the distribution of wealth may all be radically altered. If the migration is rapid or accompanied by major shifts in thought paradigms, these alterations may be traumatic or lead to conflict. If we indeed experience a gradual migration of human consciousness into the synthetic universe, we will also have a growing need for economic, political, social, and cultural expertise to deal with the difficult issues that will arise. From first-hand experience, I believe there is room for both optimism and pessimism on this score. On the positive side, I've had many inquiries from graduate students and junior professionals who are eager to become experts on synthetic worlds. On the negative side, all of these junior-level thinkers, without exception, have been hampered by an (understandable) bafflement on the part of their supervisors that something really important might be

happening in video games. An aspiring doctoral student forwarded me a thesis proposal in which the *main* research question was simply, "What are MMORPGs?" Since such a descriptive question seemed, to me, subpar for PhD-level work ("What are automobiles?"), I pressed for an explanation and was told that the student's major professor had required it. That kind of story tells me we have a long way to go before we will have a significant number of influential people who can make reasoned judgments about this technology. Indeed, part 1 of this book was, in one sense, written precisely to facilitate interactions between junior professionals and their superiors. Were I a graduate student, or a young journalist, or a middle-management executive, facing the inevitably crabby objections of the senior people ("Wait—what *are* these 'm-mor-pegs'?? And, why in the world would you think that anything in a video game could be that important??"), my answers would be summarized out of part 1.

But let's assume that this aspiration has been realized—that we know what synthetic worlds are—and that we also begin to observe a migration in progress. Part 2 considers some of the consequences. But I'd like to stress from the outset that there is more here than an effort to point at something on the Internet and say, "Look! It's amazing, new, and radical and will change the world!" While I do intend to talk about the conceivable broader implications of synthetic worlds, I don't intend this to be a book of hype. I'm not an eager technovisionary or an ambassador of digital hipness from the Millennial Generation. No, I am a middle-aged guy who works in Social Studies. My objective in part 2 is to discuss many of the hard economic, political, and security-related questions that synthetic worlds bring up.

Consider the questions about commerce and economics discussed in chapters 7 and 8. On the one hand, commercial activity seems to emerge automatically and with great gusto in these worlds. On the other, every world-builder and traveler knows that most synthetic worlds have what can only be termed "weird" economies at the macroeconomic level. Some say they are broken. Some say they work just fine. Some of these economies seem to function well, but they are no fun. Others function very poorly and are a blast to play in. The hard questions here are, What could the economy be, in a world where every physical object can be costlessly rendered in whatever quantities desired? And what *should* it be? Whatever the answers may be, it is apparent that nobody is waiting for them. Rather, users of these worlds are forming economic organizations left and right, some incorporated on Earth, most not. They mine items and gold from the worlds and sell them on eBay; they offer guide services; they speculate in rare quantities; and they make money. By some reports, some of them are making quite a bit of money. At the level of a few video games, this is at most mildly interesting. But what happens if these video games emerge as a parallel economy of significant proportions? By my own estimates, the collective volume of annual trade in synthetic worlds is, at this writing, almost certainly above $1 billion.[5] In other words, it already exceeds the total sales of a few real countries. What is true of sales is also true of total synthetic world production, as well as production per capita. Indeed, GDP per capita inside synthetic worlds is far higher than in the real world's poorer economies, such as those of India and China. There are clear implications for labor markets and globalization. Yet we have quite literally no formal expertise on how these economies are designed and managed.

Another hard question involves governance, discussed in chapter 9. Who owns synthetic worlds; who should own them; and how should relations between the owners and the residents/players be structured? And again, what's the objective? Should a synthetic world government be held to the same standards of performance as we apply to Earth governments? Or is the objective simply to have fun? Do all good governments, as defined by the great, and dead, political philosophers of the Western tradition, maximize the amount of fun to be had? It seems to me that this last question is a very hard question indeed, one that is intrinsic to the very act of world-building, and may be closer to the core of contemporary politics than we realize.

All of these questions raise another: Who (if anyone) should regulate transactions and relationships across the silicon border, and how? To what end? Who are the constituents of these policies and what are their rights? How are they represented? What rights do Earth governments have in regard to threatening organizations that exist only in synthetic worlds?

Finally, consider questions about war. As I discuss in chapters 10 and 11, people who play these games fight, constantly. Sometimes they fight in rooms that look like the room you're standing in. That's right. It would be easy for some computer whiz to build the bookstore or library (or trash can, I suppose) you're browsing through. He could put in all the shelves, indeed all the books in all the shelves. He could make the stairs and the counters and the doors and the locks on the doors. Using current AI technology, he could build robots that would do basically what you have been doing, wandering about pseudo-aimlessly, finding this section of books, glancing at various titles and then, as the result of some horrible cognitive catastrophe, making the mistake of sliding *this* book from among its neighbors, opening it, and flipping through its pages. Our coder could build the security cameras too, and the alarms, and he could endow the security guard's pistol with the correct range, line-of-sight, rounds per minute, and impact effects. Then he could make characters for himself and his friends, wrap sticks of dynamite to their chests, and place detonators and guns in their virtual hands. And then they

could all work on their timing for the Big Event. Practice, practice, practice: since it is a virtual world, they can take all the time they need, to figure out exactly where to stand to block off your exit, where to shoot to kill the security guard who rushes down the stairs, and where to shout to get the attention of the TV crews who will accumulate outside. They could do all of this in cyberspace, using existing synthetic world technology. In fact, for all we know, they've already done it. The room you are in right now may already be part of someone's Master Plan. How would you find out? How do you stop it?

In brief, this is not a book about that wacky, wonderful Internet, wacky and wonderful though it may be, but rather about some very difficult issues that may be appearing on the horizon. It's a book that tries to apply age-old reasoning from the social sciences to some age-old problems that are resurfacing, in a new way, in a very different kind of place.

One Guy's Journey into Synthetic World Expertise

I mentioned at the beginning that, as strange as some of these possibilities may seem, it is almost as strange to be claiming expertise about them. What makes a person write a book about the serious consequences of video games? The story of my own transition into this dubious status might be illuminating, I think, because it reveals how strong synthetic world technology has already become as a force for change.

You see, to me, synthetic worlds did not seem to be anything more than fun, and perhaps funny, at first. Well before I was an economist, I was a person who played lots of games. I started with geeky board games as a teenager (*Panzerblitz*, an old Avalon Hill title, was my first) and advanced to multiplayer political games (*Kingmaker*, Avalon Hill) and fantasy games (*Dungeons and Dragons*, TSR) in college. My college years were also the golden age of arcade video games, and I spent as much time as anyone shoving quarters into slots.[6] For some reason, the powers that be at Georgetown decided that a great place for the game arcade would be directly below a classroom, and I spent many class-time hours completely distracted by the bings and beeps incessantly sounding off a few feet below me. I did manage to graduate, and I went on to higher training in the field of economics. Feeling that it was important to be as morose as the discipline I had chosen, I put the games away during graduate school. At about that time, however, personal computers became a necessity for anyone doing advanced economic research, and I saw no reason why I should not take a few moments in between statistics sessions to play a computer game or two. I mean, here was all

the fun of the arcade, without spending 25 cents for each play! And no lines, either. By the end of my graduate training, I had developed a rudimentary understanding of economics but a deep mastery of the subtleties involved in stealth fighter tactics (*F-19*, Microprose), urban planning (*SimCity*, Broderbund), global domination (*Civilization*, Microprose), and the properties of odd-shaped and conditionally self-exploding blocks that fall sequentially through a low-gravity environment (*Tetris*, AcademySoft). At my first job after graduate school I remember having an uncomfortable discussion over cocktails with a very respected senior faculty member, after he revealed that he was himself a fan of the game *Civilization*. We enthusiastically traded strategies for a while, until the bemused smiles of those standing around made the whole thing feel too weird. So we went back to talking about German economic policy. It certainly occurred to no one at that party that economic policy might be happening *inside* a game within a few years.

I only discovered that in April 2001, when I started to play the game *EverQuest* (Sony Online Entertainment). This was a fantasy role-playing game, not unlike *Dungeons and Dragons*, except that it all happened online. I picked it up one evening because I thought it might be fun to act like a wizard or warrior, along with other people, in a computer game. When I first entered the game world, I was shocked at how complex the place was. Learning how to move and speak took quite a bit of time; learning my way around the huge world, with no maps, was also a challenge. Then there was the entire system of self-defense, combat, and spell-casting to be learned. Lastly, and most important, there was a rich and well-developed player society, whose language, culture, and norms were entirely foreign to me. Inside the game, there were thousands of players present at any one time, and outside of it, there were thousands of websites devoted to various aspects of gameplay. It was like no game I had ever played before. It was an intellectual challenge simply to understand *what EverQuest was*. It was a slice of practical virtual reality, as I've said, but since I had completely missed the VR hype of the previous decade, that similarity was completely lost on me. I just thought it was an incredibly interesting and deep implementation of network technology that also happened to be fun to play around in.

While having fun and trying to figure the whole thing out, however, I noticed immediately that economic transactions between players were an incredibly important part of what was happening. Within the game, in the cities, people were constantly shouting offers to buy and sell goods and services. "Want to sell Electrum Ruby Ring, only fifty platinum pieces!" "Looking for teleport to the Butcherblock Mountains, will pay." And so on. This buying and selling generated much of the sense of vitality in the place. Commerce was important outside the

game as well, revealing how many people seemed immersed in this place and what they were willing to pay, in real money, to do things there. A number of fan sites listed in-game prices for items. Moreover, online auction sites hosted US dollar markets for game goods, including game currency. At these sites, people would auction off the play money of the game for real US dollars. Intrigued, I started to surf these sites from time to time and jot down the implied exchange rates. Inside the game, I took note of the items that sold for the highest prices, and I tried to explore the world to see why these items were so expensive.

As it was, the whole thing made economic sense. The prized items made a character in the game more powerful. With the right equipment, a character could cast more powerful spells, push deeper into dangerous dungeons, and travel more quickly and safely across the continents. Since these activities were fun things to do, according to the players, it made sense that items that enhanced those activities would be highly demanded. The online trade in dollars also made economic sense. Some people who had lots of dollars happened to want more *EverQuest* play money, and some people who had lots of *EverQuest* play money happened to want more dollars. After all, you couldn't buy super magic swords using dollars; you had to trade the dollars for *EverQuest* money first, and then use the *EverQuest* money to get the sword you wanted. It made sense that the people with the excess stocks of *EverQuest* play money would be willing to exchange them for dollars, and it also made sense that people with relatively large stocks of US dollars would be willing to exchange them for *EverQuest* play money.

So, here was a great big fantasy world playing host to a basically normal economy. It occurred to me to write a tongue-in-cheek paper about that. Wouldn't that be a hoot, I thought. A paper that would look like the *World Bank Economic Report for Poland*, except it would be *Castronova's Economic Report for EverQuest.* And this would not have been such a strange thing for a midcareer economist to do. There is actually a tradition of tongue-in-cheek research in economics, in which authors will write careful papers about the economics of trivial, funny things. There is a well-known paper about the economics of the American TV game show *Jeopardy!* for example (Metrick 1995). Papers like this can be quite entertaining for economists—remember that at conferences they would typically confront such a paper following a less-than-scintillating talk on "Capital Taxation in Dynamic Models with Perfect Foresight: The Case of New Zealand." After something like that, a paper about the economics of a game show can seem very interesting indeed. I felt that a paper about the economics of *EverQuest* would have a similarly pleasing effect: a joke, but an entertaining one at least. In the dismal science, it's not too hard to be a decent clown.

Tongue firmly in cheek, I began to record all the data I could lay my hands on—whatever seemed relevant to an economic report on this odd place. I got price data from in the game and from fan websites. I counted in-game populations as many ways as I could think to do it. I surveyed players and asked them about their investments of time and money into their characters. I tried to figure out how much effort it took to produce various commodities and services. I bought and played other games, to get a sense of how *EverQuest* was unique and how it was typical. And I scanned the online auctions, to try to measure the real-world money value of all these prices and production levels. I tried to get some measure of everything in this world that would matter to a Central Statistical Office or Census Bureau.

All of this took several months in the Summer of 2001, and it had to compete with teaching, my ongoing research on social policy and the welfare state, and settling in to a new apartment with my wife. However, the *EverQuest* topic steadily occupied more and more space in my mind. Those of us who do research for a living actually don't have much control over the subjects we study; thinking goes on 24 hours a day, and it seems that if a topic so desires, it can force a higher priority for itself than we want it to have. So it was with the *EverQuest* Economic Report. I kept thinking about it all the time—while commuting, on waking up, before going to sleep, sitting in church, waiting at the doctor's office. I picked up a few books about technology, computers, and the Internet. I rented science fiction films like *Brazil* (1985) and *The Matrix* (1999). I subscribed to *Wired.*

It turned out that there were dozens of synthetic world games like *EverQuest* in development at the time, and at one point I stumbled across one called *Project Entropia* that had some extremely interesting features. This game world would allow players to buy and sell items using Earth currency, and its business plan was to make the synthetic world so large that all humans—all 6 billion of them—could be in it, all the time. The world owners hoped to entice the entire human tribe to visit the game world, and then make billions by selling advertising space and virtual storefronts within it. The whole thing sounded preposterous in the moment, but it was a rather breathtaking vision of the future all the same. It led me to ask whether such a thing was even possible.

After a bit of further reading, I became convinced that it is, indeed, possible, if not now, in a generation or two. Ray Kurzweil is a brilliant and well-respected technologist, an inventor of marvelous devices ranging from sophisticated electronic music synthesizers to reading machines for the blind. In *The Age of Spiritual Machines* (1999), Kurzweil discusses the nature of computing and its prospects for growth. He makes a number of arguments about the future. Not all of them were persuasive to me at the time, but one that seemed very hard not to believe involved

the raw amount of computing power that humans would come to possess. The short, simple point in Kurzweil's book is that humans are going to have *a lot* of computing power. This seems likely even if nothing especially exciting happens by way of technological innovation. As a result, the experience of every generation, for quite some time, will be that the next generation grows up with incredibly more powerful computers. My generation, those who played arcade video games as teenagers, is only the first of many to have this experience. And it follows that almost any computational problem that can be imagined by a single person right now will eventually be mastered by a computer. The way Kurzweil expressed it, at some point in the twenty-first century the standard personal computing device will have as much computing power as a single human brain. Not long after that, it will have the computing power of *all the human brains that have ever lived.*

One important use of all that power, according to Kurzweil, will be to upload brains and recreate consciousness inside silicon. I doubted that on reading it, but I was not following that line of argument very closely anyway because I was thinking about *EverQuest.* Add computing power to a game world and you get a place that's much bigger, much richer, and much more immersive. The robots running around in it, humanoid and unhumanoid, are smarter and act more and more like real people and real monsters (if there is such a thing). Add *immense* computing power to a game and you might get an incredibly realistic extension of Earth's reality itself. The place that I call "game world" today may develop into much more than a game in the near future. It may become just another place for the mind to be, a new and different Earth.

It turned out that Kurzweil was not the only one making such judgments about the potential computing power of humanity. As I mentioned earlier, a number of authors have made similarly persuasive arguments (Moravec 1999; Joy 2000). While some have argued that increases in raw processing power do not automatically convert themselves into improvements in economy and society, not too many dispute that increases will happen.[7] It's hard to see how they might not happen. Any piece of matter that can be switched between two positions can be used to store binary data and process commands on it. We know we can compute on silicon, a material whose raw computational power doubles every two years or so. But we can also compute with other materials, such as DNA or quantum particles. We have, already in place, extremely strong economic incentives to encourage researchers to discover these new computing methods. Thus, from an economic standpoint, it seems quite likely that if there are sources of computing power in the materials around us, they will be tapped. The extremely rapid growth we have observed so far—from glowing green text to immersive 3D forests in just *twenty years*—seems likely to continue, whether on silicon or something else.

The prospect of continued growth in computing power had a serious implication for my research into *EverQuest.* What was just a game world in 2001 might be a richly simulated and completely immersive mental environment in 2101. Perhaps by 2101, it would be trivial to simulate the entire scope of known human history from 10,000 BC to AD 2000 at the level of individual synapses of all the brains that lived in that period. Or perhaps worlds like that would only be possible in 2200, or 2300. And perhaps no one would want to simulate that history; perhaps we would all prefer a different history. While we are at it, why not just simulate our future instead? Those pesky stars and galaxies are so far away. Under Einstein's rules we will probably never get there. But if we simulate a space-time world that allows us to travel there in our minds, if not our bodies, we can go there right away. Whatever these simulations will be like, and at whatever pace they appear over the next few centuries, it seems clear that they will be an important host of human mental activity. How exciting, then, that early outlines of their characteristics can be seen today, right now, by looking at their ancestors: game worlds like *EverQuest.* My silly paper about a computer game thus turned into a more general study of the economic and social life just starting to develop inside this silicon shell.

Because it had to deal with such an unusual and emergent phenomenon, the paper that I eventually wrote ("Virtual Worlds: A First-Hand Account of Market and Society on the Cyberian Frontier") ended up looking nothing at all like a traditional economics paper. It had a wealth of economic and social data about the world of *EverQuest.* I found, for example, that you could work in *EverQuest* and earn something like 300 platinum pieces (*EverQuest*'s currency) an hour, on average. If you converted that into real money by selling it online, it would come to about $3.50. So you could earn a poverty-level wage, in real dollars, by "working" in the game. I also found that the ongoing exchange rate between platinum pieces and dollars was higher than that of the Japanese yen, the Korean won, and (at the time) the Italian lira. When I looked at population numbers and estimated the market value of production on an annual basis, that came to over $2,000 per person. In other words, I had found that the Gross National Product of Norrath, per capita, was about the same as Bulgaria's and four times higher than China's or India's. This game world, though small, was host to quite a bit of productive activity.

Aside from economic data, the paper was also filled with vignettes and stories about the strangeness of life in Norrath. Some kind of storytelling was essential, because the world being described was foreign to the general economist reader. This was almost more deep-immersion anthropology than economics, but, frankly, economists generally don't believe in deep-immersion research and they certainly don't care too much for "vignettes." As a result, I was

not sure about the prospects for publication. I had a research affiliation at the CESifo institute at the University of Munich, however, that would allow me to place up to four draft papers annually in their working papers series. CESifo Working Papers could be downloaded through the Social Science Research Network, the main online clearinghouse for research in law, management, economics, and finance. So I decided that I could just release the paper that way. I could find out whether anyone had an interest in an unusual economics paper about an unusual thing, a thing that might be serious or might not be, depending on how you feel about the future of computing and the Internet. The paper became a CESifo Working Paper in December 2001 and was listed at SSRN.com in January 2002.

Now, at that time, economics was not the most popular thing being downloaded on the Internet, not even at SSRN. Download a finance paper, and you might learn how to make some money in stocks. Got a legal issue? Download a law paper. Downloading an economics paper was typically less likely to get you anything immediately useful. As a result, the top downloads at SSRN were all law, finance, and management papers that had been around since the inception of the site (1997) and had accumulated thousands of downloads. Comparatively, a typical SSRN paper in economics might have gotten one or two downloads in its entire lifetime; a popular paper might have had one or two a month; a paper assigned by a professor to his students might get 50 or 100 downloads. I thought "Virtual Worlds," as a quirky paper, might get 100 or so downloads. Maybe some econ grad students might use it to have something edgy to talk about over lunch.

Given these expectations, I was pretty shocked when the paper had 500 downloads the first day or so. My wife and I went out for a nice meal. When it hit 1,000 downloads a few days later, we went out for another nice meal. At 5,000 downloads a couple weeks later, we went out for still another nice meal. At 10,000 downloads, about two months later, we went out again, but agreed that our budget and waistlines couldn't handle any more nice meals. But people kept downloading the paper. By summer 2002, "Virtual Worlds" was pushing 20,000 downloads and had entered the all-time top-ten download list for SSRN, the only paper published since 1997 to have done so. It was also the all-time number one download for economics papers. Well, I couldn't avoid the conclusion that the paper had made some kind of an impact with lots of people, but, for the most part, I had no idea who these people were or why they thought the paper was worth downloading.

Then the email started pouring in, and the phone started to ring. Game companies wanted advice on developing in-game economies. Mainstream journalists wanted interviews about these insane people paying real money for magic swords. Business and tech journalists wanted to see if the numbers really added up and whether the person behind them was really an economist or just some charlatan. Entrepreneurs of various stripes wanted ideas for making fortunes. Academics in the arts and humanities wanted me to speak at conferences about digital culture. Government agencies wanted info on how synthetic worlds could be used for training and policy analysis. Intense game players wanted me to study the economy of the game *they* really thought was cool. Students at all levels wanted to know how to get gaming into their research, or how to get jobs in the gaming industry. Lawyers wondered about property issues. I learned that this exploration of a synthetic world was apparently a very serious thing indeed, to people operating in an incredibly wide spectrum of human affairs.

It was funny that the only people who did *not* call were other academic economists. Years later, I did begin to have conversations with prestigious economists that began "So *you're* Castronova! My graduate students have been citing this thing about video games and I keep wondering, who the hell *is* this guy??" But in the first months of my paper's journey around the Internet, it was clear that the profession judged it to be hype, or some kind of fad, and steered away. So, as I feared, my efforts to place "Virtual Worlds" at traditional journals were turned down. "Not the right kind of paper for our journal," one editor said. "The subject is not ripe," said another. Commenting on a later paper on the same subject, a reviewer for the prestigious *American Economic Review* opined: "Basically, these [results] tells us [*sic*] something about the distribution of tastes of a small group of individuals who play virtual reality games. Personally, I'm much less interested in the pricing of the characteristics of "virtual" things than of those that are real" (quotation marks in the original).[8]

I decided that since I had enough normal research to keep my career going, there was no need to formally publish the *EverQuest* piece. It would be fine if it remained just a Working Paper. It certainly had had more impact in that form than any journal article I had ever written.

The experience left me at a crossroads in terms of my career, however. Evidently, this was not a line of research that was going to be well received by the hierarchy of academic economists. On the other hand, it *was* a line of research that seemed to contribute something tangible to the thinking and working of a great many people. Now, according to one view of academic freedom, it is given to academics so that we can explore things that need to be explored, regardless of the consequences; that's the duty that goes along with the benefit. I never expected to be someone who actually had to rely on the academic freedom argument to make a decision about my career, but in the end, that's how it was. I had important people from many different areas urging me to pursue the study of this

topic in one direction or another. My paper had raised many more questions than it answered. I decided to keep working on some of these questions, and that's why I am writing this book.

But it is important to stress the moral of the story: No one had to be persuaded that synthetic worlds were important, they made the case themselves. I just wrote down a description of what was happening inside a video game, and how the technology behind it was being considered by others for much more expansive uses. The paper was never approved by a board of editors and still has not appeared in print. It was just a document floating around the Internet, a free self-publish by an unknown economist from an obscure teaching college outside Los Angeles. But its subject was powerful enough to stir up interest in places like the *New York Times*, the Pentagon, and the grimy offices of graduate students and game coders alike. I can't possibly take credit for generating all that interest, even granting myself a megalomaniacally generous assessment of my own talent as a writer. No. The paper got passed around because many people sense that something really is happening, out there, on the other side of the screen.

An Overview of the Contents

The subject of the book is a generic "synthetic world," by which I mean any computer-generated physical space, represented graphically in three dimensions, that can be experienced by many people at once. It makes no sense to list the specific games and worlds that fit in this category at the moment; this phenomenon is moving so quickly that any list would be out of date by the time the book hit the street. To give an idea, however, I take "synthetic world" to include all the MMORPG worlds in the game industry. These are role-playing games about dragons and spaceships, and are exemplified by the game *EverQuest*. I will also include multiuser social worlds, such as *The Sims Online* (Electronic Arts), that typically do not have an explicit gaming, combat, or competitive aspect. At times, I will discuss multiplayer combat games such as *Half-Life* (Sierra), although players of these games cannot have the same kind of persistent presence as they do in MMORPGs and social worlds (for reasons to be discussed later). I will not include such things as text-based chat rooms or MUDs (Multi-User Dungeons), or online games such as online poker or online checkers; these technologies do not depend on a 3D, world-like physical interface. Of course, no single world sufficiently describes the entire universe of worlds now emerging; therefore, the book will have to focus on what seem to be generic and persistent features of all the worlds as a group.

Just what is generic and persistent about synthetic worlds, and what is not, will require some guesswork about future developments. For example, at the moment it is fairly easy in many Internet communities for users to remain anonymous, and thereby to be shielded from the consequences of their actions. I believe this is a temporary situation. Human societies rely so much on reputation for their basic functioning that online anonymity seems unlikely to persist in any significant way. Similarly, there are a number of technological issues that affect life in synthetic worlds at the moment, that are more likely to be solved rather than accepted. Right now, for example, it is hard to get a mob of 50 or so people together at one time at one place online, because the required digital video processing often becomes more than a contemporary system can handle. Nonetheless, I will not talk much about the difficulty of forming armies or markets in synthetic worlds. On the contrary, I will be assuming that armies and markets will be readily available forms of social organization; they are so important for human collectivities that the coders will undoubtedly find some kind of fix or workaround using the computation resources to come (indeed, to some extent, they already have). While assuming certain technological and social advances can be risky for a book such as this, it is warranted in some cases. A book written in 1920 about the emerging film industry might have more or less safely assumed that both sound and color would become the norm eventually.

The book has three major parts. As mentioned above, the first part is devoted to an overview of what synthetic worlds are, how they operate, and how they interact with the world of Earth. Chapter 1 gives a description, based on my own experiences in dozens of these games, of what it is like to spend time in a current-generation MMORPG. In chapter 2 I'll give an overview of the people who go to these places, their numbers, and why they make this choice. Chapter 3 sketches the technology behind these world simulations; here, I make some more or less well-grounded predictions about the near-term technological potential of these worlds. Chapter 4 describes the "institutions"—the patterns of culture and behavior—that these places induce. Chapter 5 concludes the first part by giving a brief overview of the industry that makes these worlds, its corporate structure and atmosphere, and the likely equilibrium state of the market as it evolves.

Part 2 then considers the potential effect of synthetic worlds with respect to a number of important issues. Chapter 6 describes many of the boundary-violating practices that have become common in MMORPGs, showing how fuzzy the line between the game-space and real-space has already become. Chapters 7 to 11 then discuss some of the implications of these fuzzy boundaries for economics, politics, and security.

Part 3 wraps up the book with two chapters, one an assessment of policy issues, and another that attempts to make longer-run projections about the ultimate meaning of this technology.

Indeed, what are the right policies toward these new places? This is a difficult question; we are not sure what these worlds will be, let alone what they should be. People in different positions have different opinions. Those who own worlds and want revenues from them have a certain set of goals in mind. Those who are heavily invested as users of a world will have a different set of goals. Users who have spent much time in one role in a world (as a warrior, a trader, a lover, or what have you) will not have the same idea of "ideal" policies as those who have specialized in a different role. Not to mention those who have not entered the world yet. And we can of course conceive of a common good here; there may be some policies that are commendable as is, either because they offer the best balance of conflicting interests, or because they are simply the right thing to do. While I would like to claim neutrality with respect to the various interests at play, that is going to be impossible in practice. I am going to be arguing, at various points in this book, that it's better to do things a certain way. When that happens, I will be trying to side with the common good as much as possible.

My conception of the common good is fairly simple: it's the collective well-being of the users. And by "well-being," I do not mean a strictly utilitarian conception. The mere fact that one synthetic world sells better than any other is not, in my view, conclusive evidence that it is a better world than the others. It is strong and persuasive evidence—to assert otherwise is elitist and tyrannical—but it is not conclusive evidence, and the reason it is not conclusive involves the nature of free will. Frankly, I fear the power of a simulated world to alter the mind's conception of its own desires. Many people seem to become heavily invested emotionally in the rather crude synthetic worlds we already have, and some spend almost every waking hour there. Is that the result of a rational choice, or rather of some form of chemical response treadmill similar to nicotine? Many of these worlds are designed to utterly and completely enclose the user's consciousness in an envelope of the coder's design.

What are the long-run implications of that practice? Quite a few social scientists, especially the economists who trained me, believe that tastes are fundamental. But no economist denies that, in the very long run, the social and physical environment has a significant effect on what we like and don't like. Saharan nomads don't tolerate cold weather as well as the Inuit. With synthetic worlds, we are talking about a technology that can dictate the physical environment (and to a great extent, implicitly, the social environment) within which a mind matures, for months or years or decades. The preferences that economists and other

rational choice social scientists treat as fundamental and unchanging are very much changeable here. That puts them in play, and forces us to think about the Good and Bad of a world in a way that is not entirely dependent on the tastes of the people in it.

There's nothing radical in this. For decades, economists such as Nobel Laureate Gary Becker (1996) have theorized about the way preferences may change. Libertarian theorist Robert Nozick (1974) has no trouble rejecting an existence based on permanent immersion in a positive experience-producing machine: such a life would be horrible, even if the immersed person said it was wonderful. Because everything about a synthetic world can be revised by the coding authority, the world needs to be Good as well as popular. We cannot simply assume that a world that commands the reverence of its people is necessarily the best world for them.

Ideally, the synthetic worlds we revere and feel most comfortable in will also be Good worlds. We will like them, *and* they will enhance human dignity and well-being. That is my hope, anyway. The Earth is very nice, but there are experiences we can imagine in our minds that we *cannot have here.* We poor Earthlings simply cannot explore the surface of Pluto. We really cannot magically heal our best friend's sickness with a touch of our hand. We cannot cast a spell and fly away. We cannot switch from male to female and back again; we cannot become fat or thin as we wish; we cannot have the strength of a bear or the speed of a rabbit. We cannot directly observe the signing of the Magna Carta, Pickett's Charge, the remilitarization of the Rhineland, D-Day, the Defenestration of Prague, the severing of Jenkin's Ear, or the invention of Silly Putty. And certainly, we cannot cast immense fireballs that immolate things that annoy us. Much as we might want to do that, it remains, for the known future, something not possible for ordinary humans like you and me. On Earth, anyway.

Looking beyond these simple joys of immersive, interactive entertainment, however, it should be stressed that synthetic worlds may eventually make contributions to human well-being that will be judged as extraordinarily significant. We live in a mobile world, and family ties often fall victim to the locational demands of the marketplace. Synthetic worlds may become a place where families can maintain their togetherness, even though the members live thousands of miles apart. Moreover, synthetic worlds may allow us to experience human social life in an environment in which many characteristics of the body are no longer fixed endowments but have become chosen attributes. People entering a synthetic world can have, in principle, any kind of body they desire. At a stroke, this feature of synthetic worlds removes from the social calculus all the unfortunate effects that derive from the body. Imagine a world in which all aspects of our physical

appearance were under our control, so that all variations in thin, heavy, tall, small, dark, and light were all voluntary. We are poorer for being still unable to experience such a world.

To sum up: synthetic worlds hold immense promise as places where humans can enhance their Earth experiences with ones drawn directly from our glorious collective history and imagination, all without bearing some of the burdens that adhere to the Earth bodies we were born with. Ensuring that the technology serves such a marvelous end, rather than a less happy one, is the real challenge for the next few decades. We will be less likely to meet that challenge the longer we treat video games as mere child's play. Multiplayer online video games—avatar-mediated communications systems, in essence—have become too fraught with heady implications to be ignored any longer.

PART I

THE SYNTHETIC WORLD: A TOUR

IF YOU ARE NOT AN AVID PLAYER OF VIDEO GAMES AND YOU DO not spend very much time online, your first questions about these places will not be about implications but about basic features. What phenomena are being considered when someone uses the term *synthetic worlds* (or the more popular synonym, *virtual worlds*)? The chapters in part 1 attempt to describe common features of all the synthetic worlds now in existence and likely to come online in the next decade or so. Chapters 2–5 talk about the users, the technology, the behavioral pattern, and the market structures that most synthetic worlds seem to share. To begin, however, the first chapter takes the reader on a walk-through of a typical synthetic world. How is it that you "go" there? Why does it feel any different from, say, opening an email program or a game of computer solitaire? Once you are there, who else will you find and what will you do with them?

If you're already quite familiar with synthetic worlds, you could skip the walkthrough and go on to chapter 2. On the other hand, even the most experienced gamer might find it interesting to see the level at which I've chosen to describe a typical day in a MMORPG. Much that seems second nature to the experienced gamer is, in fact, a thoroughly befuddling recoding of expectations for the rest of us. It's in the interests of everyone—those who know the technology and culture inside and out as well as those who've never heard of these things before—to realize how much experiential distance separates them.