

Castronova, Edward

Working Paper

Theory of the Avatar

CESifo Working Paper, No. 863

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Castronova, Edward (2003) : Theory of the Avatar, CESifo Working Paper, No. 863, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/76336>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THEORY OF THE AVATAR

EDWARD CASTRONOVA

CESIFO WORKING PAPER NO. 863

CATEGORY 2: PUBLIC CHOICE

FEBRUARY 2003

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the CESifo website:* www.CESifo.de

THEORY OF THE AVATAR

Abstract

The internet has given birth to an expanding number of shared virtual reality spaces, with a collective population well into the millions. These virtual worlds exhibit most of the traits we associate with the Earth world: economic transactions, interpersonal relationships, organic political institutions, and so on. A human being experiences these worlds through an avatar, which is the representation of the self in a given physical medium. Most worlds allow an agent to choose what kind of avatar she or he will inhabit, allowing a person with any kind of Earth body to inhabit a completely different body in the virtual world. The emergence of avatar-mediated living raises both positive and normative questions. This paper explores several choice models involving avatars. Analysis of these models suggests that the emergence of avatar-mediated life may increase aggregate human well-being, while decreasing its cross-sectional variance. These efficiency and equity effects are contingent on the maintenance and protection of certain rights, however, including the right of agents to free movement, unbiased information, and political participation.

Keywords: information and internet services, computer software, equity, justice, inequality.

JEL Code: L86, D63.

Edward Castronova
Cal State Fullerton
Department of Economics
Fullerton, CA 92834
U.S.A.
ecastronova@fullerton.edu

I. Introduction

In 1996, 3DO studios released a computer game called “Meridian 59.” The game allowed players to enter the gaming space by internet connection, and it allowed a large number of them to enter at the same time. Each player’s screen became a viewport on this world; looking at it, one could see houses, building, fields, and other players. The players appeared in the game world as graphical objects that looked like human bodies. Like human bodies, these graphical objects had to run across space to get from point A to point B; they had to climb stairs or ladders to get on the roof of a building; they had to watch out for dangerous monsters, or, more precisely, computer-driven graphical objects that looked like monsters, whose primitive artificial intelligence code instructed them to kill anything that looked human. To avoid the monsters, or kill them first, the players had to give their human representations similar kinds of orders: look over here for the monster; run to it; get out a sword; attack. Because there were other players in the same playspace, players could access more complex commands: ask your friend to come over here; agree to attack the monster together; split up loot from the monster; trade some of the money to a third player for a magic helmet. Tens of thousands of people came to the world of Meridian 59, the world’s first large scale 3D first-person multiplayer game. It still exists, as of this writing, in pirated copies; people still visit from time to time.

Meridian 59 was the result of a rapid development of shared virtual reality spaces that first began to emerge on networked computers in the late 1970s. By 2001, the internet had spawned more than fifty such spaces, some nothing more than graphical chat rooms, others consisting of huge worlds with complex economies, societies, and polities. Games like EverQuest from the U.S. and Lineage from Korea came to be populated

globally by tens of thousands of people at a time, 24 hours a day, year round. A not insignificant fraction of the players of these games reported that they felt themselves to be not game players, but citizens of a new world (Castronova, 2001). For many, life in the virtual reality spaces seemed preferable to life on Earth.

Meanwhile, the industries that produce shared virtual reality environments have continued to develop rapidly. Advances in connectivity (bandwidth) and interfaces (haptic devices, heads-up displays) have been driven by technology. The amount of available content (narratives, folklores, backstories) has expanded; the video game industry has begun to surpass the motion picture industry in gross revenues.¹ Large corporations have begun to transform deeply enmeshed cultural icons into virtual reality spaces. Sony maintains a world of Star Wars Jedi Knights; Sony, Microsoft, Electronic Arts, and others compete in the space of Tolkienesque worlds of elves and hobbits; EA supports a vastly popular virtual world of Sims, beings who would be unremarkable except for the fact that they exist only in a cyberspace reality. As the phenomenon continues to grow, the aggregate amount of time devoted to shared virtual reality spaces seems likely to rise from today's tens of thousands of person-years into the hundreds of thousands or perhaps millions.

Devoting time to a virtual reality space is a choice that, while unusual in some respects, is amenable to economic modeling. This paper presents simple models of the choices involved in visiting virtual worlds. These choices include how much time to spend, but also where to travel, and, especially *how*. The latter choice emerges because virtual reality spaces offer visitors the opportunity to take on a physical self-representation that is typically different from the body they have on earth. The physical

representation of the self in virtual reality – the **avatar** – is an important aspect of the choice problem, as we would expect that people would gravitate toward those worlds that offer them their ideal avatars.² In that sense, an economic model of participation in virtual reality reduces to a theory of avatar choice. A theory of the avatar can make both positive and normative statements, positive in predicting who will choose what avatars and how much time they will spend in them, and normative in describing the welfare consequences of an avatar-mediated life.

The paper is organized into five sections. Section II presents a simple time-allocation model of avatar demand, while Section III extends that model to consider the choice of avatar attributes (height, weight, hair color, etc.). Section IV discusses avatar choice as an interactive game among many agents, and considers the possibility of engineering these interactions to produce more equality. Section V explores an example in the context of team-based production, by developing a simple general equilibrium model of output and avatar choice. Section VI discusses the potential anonymity that avatars provide and the potential effect on social norm and reputation models. Section VII concludes with a discussion of policy issues raised by the emergence of the avatar.

II. The Avatar

As defined above, the avatar is the representation of the self in a given physical environment. The Earth has a physical environment, with certain laws of motion, gravity, force, and so on. Things that happen on Earth are seen, heard, and felt by us, through the medium of our physical senses. Thus, when our minds experience the Earth, they do so through our bodies. Our bodies must react to the forces imposed on them by the Earth's

environment, and when our senses detect an opportunity to meet a goal, we must direct our bodies to act in the Earth's environment to achieve the goal. Our real bodies are, in some sense, our Earth avatars: when we are in Earth, our selves are present in and represented by a body that exists in Earth, and only there.

When we visit a virtual world, we do so by inhabiting a body that exists there, and only there. The virtual body, like the Earth body, is an avatar. When visiting a virtual world, one treats the avatar in that world like a vehicle of the self, a car that your mind is driving. You "get in," look out the window through your virtual eyes, and then drive around by making your virtual body move. The avatar mediates our self in the virtual world: we inhabit it; we drive it; we receive all of our sensory information about the world from its standpoint.

Current technology does not permit complete immersion in a computer-mediated reality, of course.³ Nonetheless, even with current technology, participants in virtual worlds seem to experience those worlds fully, in the sense that they immerse their minds in the virtual place. The degree of immersion is sufficient that the player is essentially not aware of her Earth surroundings. After "jacking in," to use William Gibson's term, the agent gives primary mental attention to signals from virtual reality, with only secondary input from Earth. The player no longer seems to be "here," but rather is "there," even to the point that events "there" take on more emotional meaning than anything the person experiences on Earth.⁴ And when signals are being received from cyberspace, they come to the mind through the eyes and ears (and, in the future, nose, skin, and tongue) of the avatar.⁵ Therefore it seems reasonable to assume that the problem for the choosing agent is allocate mental attention, in units of time, to sensory input from different avatars. The

choice problem is to allocate time to avatars: For the next two hours, will I be Edward Castronova, on Earth? Or will I be Comolan Canisa, in “Norrath,” a fantasy world?

The choice will be affected by the different abilities an avatar can possess, many of which can themselves be chosen by the agent. For the moment, however, let us set aside the possibility that avatar attributes may be chosen, and concentrate instead on the basic time-allocation problem. Assume that all avatar abilities are fixed, and an agent with a fixed time budget must choose which avatars to inhabit, and for how long. Furthermore, we will assume that the mind can experience the external world only through one avatar at a time. Therefore the choice problem is to allocate time to discrete avatars.

Let the set A consist of N different avatars, indexed $a = 1, \dots, N$. Without loss of generality, let the avatar indexed $a = 1$ be the agent’s “Earth avatar.” The agent has a time budget of T hours to spend, and allocates t_a hours to avatar a . The choice problem is

$$\text{Max } U(t_1, t_2, \dots, t_N) \text{ subject to } t_1 + t_2 + \dots + t_N = T$$

where $U(\dots)$ is a continuous, twice-differentiable function, with $U_i > 0$, $U_{ii} < 0$, $i = 1, \dots, N$. The agent will allocate time to each avatar, depending on the weight that the utility function applies to the experiences the agent has when inhabiting that avatar. For some, the weight assigned to t_1 will be so heavy that the optimal allocation would be $t_1 = T$. For others, the utility weights will imply an optimal allocation with $t_1 < T$. If, for example, variety really is the spice of life, some agents will become bored with being mediated through any one avatar all the time. For such agents, the ideal distribution of time involves spending a positive amount of time in a number of different avatars.

It is probably reasonable to assume that no one would choose to devote zero hours to the Earth avatar, but it is possible. Because the physical body could be supported indefinitely, as is done in long term care facilities, people could, in principle, live their entire mental lives through a computer-mediated avatar. Such an extreme structure of preferences seems improbable, of course, but the opposite extreme may be equally improbable. Perhaps just as few people would want to live their entire mental lives mediated through their Earth avatar. Given the number of hours that are already devoted to virtual worlds, it seems that a significant fraction of the population finds a life mediated through one Earth avatar less fulfilling than life mediated through an Earth avatar and one or more others.

There is an immediate and fairly trivial implication for well-being. If, in fact, many people do enjoy splitting their time among several avatars, then the pre-computer world, in which only one avatar was available to each person, was inferior on welfare grounds. The growth of the internet, and of shared virtual reality spaces within it, has enabled new choices in terms of what kind of physical beings we inhabit. Most people have not changed their time allocation in response to this development; their utility level of $U(T, 0, \dots, 0)$ remains the same and they are indifferent to the change. Some, however, have changed their allocation from $\{T, 0, \dots, 0\}$ to $\{t_1, t_2, \dots, t_N\}$, with $t_1 < T$ and $t_a > 0$ for some a other than 1. For these agents, $U(t_1, t_2, \dots, t_N) > U(T, 0, \dots, 0)$; their well-being has risen. Thus, the advent of shared virtual reality has had a net positive effect on aggregate well-being.

III. The Avatar and the Virtual World

The preceding analysis assumes only that there is an entity, the avatar, and that people receive more or less utility from devoting time to it. We now begin to flesh out what an avatar is, and the way in which “inhabiting an avatar” can contribute to well-being. We begin by recognizing that an avatar exists in a specific world, and that both the avatar and the world have distinct characteristics that can affect the utility of the choosing agent. At the same time, the agent has certain features, inherent in her own mind, that do not change regardless of the avatar she inhabits.

Let the set R consist of all the worlds in existence, and let the vector \underline{z}_w indicate the characteristics of world w , $w = 1, \dots, W$. An element of \underline{z}_w is a variable measuring some aspect of the nature of world w . It would indicate such things as the size of trees, the number of dwellings, or the placement of cities.⁶

The agent has the opportunity to visit world w by inhabiting an avatar. Indeed, she may inhabit the same world through a number of different avatars, albeit only one avatar at a time.⁷ Let A_w be the set of all avatars that exist in world w , let there be N_w such avatars, indexed $a = 1, \dots, N_w$. Without loss of generality, assume that $w = 1$ indicates the world of Earth, and that the set A_1 consists of a single avatar, namely, our body.

An avatar possesses a unique vector of characteristics, denoted \underline{v}_{aw} , which is drawn from the set of all possible vectors of characteristics of avatars in this world, V_w . The index a indicates the avatar, while the index w indicates the world. An element of the vector \underline{v}_{aw} could indicate, for example, the avatar’s height.

The choosing agent herself also possesses non-physical attributes, denoted \underline{x} , drawn from X , the set of all possible vectors of non-physical characteristics. An element of \underline{x} , for example, might indicate the agent's IQ score.

At any one time, the agent's utility will depend on her non-physical characteristics, the characteristics of the avatar she inhabits at the moment, and the characteristics of the world in which that avatar exists. All of these characteristics affect the kind of experiences she has, which in turn affects the nature of sensory impulses reaching the mind. Some sensory impulses (seeing a beautiful landscape) will improve emotional satisfaction, while others (losing a prized item) will reduce it. The agent's objective is to manage these characteristics, so as to influence the sensory information being transmitted to the mind.

In managing these characteristics, the agent must treat the non-physical characteristics \underline{x} as fixed, because they are the same no matter what avatar the agent inhabits. However, attributes relating to the physical aspects of the agent, her body and the world her body exists in, are both open to change. By choosing a different avatar in a given world, the agent can change avatar attributes \underline{y} ; by choosing a different avatar in a different world, she can change both avatar attributes \underline{a} as well as world attributes \underline{z} . The choice problem is to determine which mix of world and avatar characteristics, combined with each other and with the inherent characteristics \underline{x} , produce the greatest emotional satisfaction.

For example, let there be two worlds, world 1 and world 2. World 1 is the Earth, while world 2 is Norrath, a shared virtual reality space on the internet. Earth consists of five continents, and an unpopulated moon. Norrath consists of five continents, and a

populated moon. In Earth, you travel from point A to point B by walking or taking a vehicle. In Norrath, you may travel by such means if you wish, but you may also teleport yourself instantly - to the moon if desired - under certain circumstances. These features of the two worlds are captured by differences between \underline{z}_1 and \underline{z}_2 . The choosing agent may have only one avatar in Earth; assume that that avatar is four feet tall. In Norrath, the agent may select from, let us assume, 8 avatars, some four feet tall, others over twelve feet tall. These features of the avatars are captured by differences between \underline{y}_{11} , the characteristics of Avatar 1 in World 1 (the earth body), and $\underline{y}_{21}, \underline{y}_{22}, \dots, \underline{y}_{28}$, which are the different characteristics of Avatars 1 through 8 in World 2, Norrath. Finally, the agent herself has an intelligence level, given by her IQ score, which is among her individual characteristics as captured in \underline{x} . The agent's decision problem is to allocate time between her avatars in Norrath and her avatar in Earth.

The utility consequences of "inhabiting an avatar," then, stem from the fact that some people might be happier living with attribute set $\{\underline{x}; \underline{y}_{11}\}$ in a world characterized by \underline{z}_1 , than with attribute set $\{\underline{x}; \underline{y}_{21}\}$, in a world characterized by \underline{z}_2 . Perhaps the agent is happiest when experiencing the world as a tall ogre than as a short human.⁸

At the most general level, a given characteristic has an impact on behavior because it affects the desirability of a unit of time spent in a given avatar in a given world. We would therefore expect the agent, avatar, and world characteristics to enter the utility function as weights, or loadings, on the units of time. In other words, the agent's utility of time spent in avatar 5 in World 7 will be weighted according to how much satisfaction she experiences when driving that avatar. The satisfaction weights on time in an avatar may be subject to complex interactions. Example: Suppose the agent likes being

a tall person, and suppose that two avatars both have the attribute {ten feet tall}. The avatar that exists in a world with characteristic {most people are 3 feet tall} will provide her with more satisfaction than the avatar in a world with characteristic {most people are 30 feet tall}. In the end, the utility of a unit of time spent in a given avatar will be a function of the full set of characteristics $\{\underline{x}, \underline{y}, \underline{z}\}$ that determine the agent's experiences.

The utility function remains defined over units of time spent in various avatars, but the attributes of the agent, the avatars, and the worlds are inserted as parameters of the utility function. A completely general statement of the utility function would be:

$$U(t_{11}, t_{21}, \dots, t_{N_1}, \dots, t_{aw}, \dots, t_{N_w}; \underline{x}, \underline{y}_{11}, \underline{y}_{21}, \dots, \underline{y}_{N_1}, \dots, \underline{y}_{aw}, \dots, \underline{y}_{N_w}, \underline{z}_1, \underline{z}_2, \dots, \underline{z}_W)$$

The direct arguments of the utility function are the time inputs to each avatar, where t_{aw} is time spent in avatar a in world w . There are N_w avatars in world w , and W worlds in total. Parameters of the utility function are indicated after the semicolon. They include the agent's characteristics, \underline{x} , the attributes of all the avatars (where \underline{y}_{aw} indicates the attributes of avatar a in world w), and the attributes of the worlds (where \underline{z}_w indicates the attributes of world w).

To simplify the problem, let the utility function be additively separable and take the following form:

$$U(\dots) = \sum_w \sum_a u(t_{aw}; h(\underline{x}, \underline{y}_{aw}, \underline{z}_w))$$

where $h(\underline{x}, \underline{y}_{aw}, \underline{z}_w)$ is function from the space of all the attributes (agent, avatar, and world) to a real number. Separability of the utility function implies that the satisfaction one receives from inhabiting a given avatar is not directly affected by the attributes of other avatars and other worlds. An example of a simple utility function of this form would be

$$U(t_{11}, t_{21}, t_{12}) = h(\underline{x}, \underline{y}_{11}, z_1) \ln(t_{11}) + h(\underline{x}, \underline{y}_{21}, z_1) \ln(t_{21}) + h(\underline{x}, \underline{y}_{12}, z_2) \ln(t_{12})$$

This utility functions expresses the utility available from three avatars: Avatars 1 and 2 available in World 1, and an Avatar 1 available in World 2. The function has Cobb-Douglas form, with the weights on the elements of utility being given by the attribute weighting functions $h(\cdot)$. The Cobb-Douglas form further requires that the weights add to one, so they would have to be constructed or normed correspondingly.

Expressing the choice problem in this form has the benefit of allowing us to see that it is the characteristics of avatars and worlds that ultimately determine the demand for time in them. When a world designer gives a certain avatar the ability to fly (an enjoyable experience for most, we assume), she increases value of the weighting function $h(\cdot)$. As a result, the demand for time in that avatar will increase. On the other hand, the ability to fly may not always be pleasing; it depends on the agent's inherent characteristics as well as the characteristics of the world itself. As a result, the weighting function $h(\cdot)$ will have to be able to express the complex interdependencies among the attributes. For example, because the emotional impact of a given attribute depends on other attributes, the weighting function will generally be non-separable. And because one can have too much of a good thing, as well as too little of a bad thing, the weighting function will generally be non-monotonic in all its arguments. Those who seek to raise demand for avatars or worlds that they have created must pay attention to these complexities, and design sets of attributes that have the joint effect of raising the emotional satisfaction of agents.

While the task of designing satisfying avatars and worlds may be complex, the fact that these worlds are in demand suggests that they will have the net effect of raising

aggregate well-being. The argument was presented in the preceding section: Before the advent of virtual worlds, each human person had exactly one avatar, her body, and she maneuvered it through exactly one world, Earth. When a virtual world comes into being, most people ignore it, and their well-being is unchanged. There are those, however, who choose to spend some time in the new virtual space. These individuals experience an increase in well-being. On the whole, then, aggregate well-being increases. There is an added nuance to the argument here, in that we can now see what causes the increase in well-being: an expansion in the portfolio of personal and environmental attributes by which a person can experience life. In some cases, a person may experience a personal characteristic as fundamentally unpleasant and undesired. Consider, for example, someone who is unable to walk with her physical body. It is easy to see that the appearance of a world in which she *can* walk would raise her well-being. But the normative argument is not about improvements in attributes alone, it is also about having a wider variety of attributes. A very tall person may enjoy being tall, for the most part, but also have great fun as a very short person, if only on occasion. Just as the introduction of new kinds of food almost assuredly raises the well-being of those who eat, the introduction of new avatars and new virtual reality environments almost assuredly raises the well-being of those who travel from body to body, and from world to world.

In sum, then, the willingness to pay for participation in a shared virtual reality environment depends on the emotional experiences that the environment provides, which in turn is a function of the attributes a participant is allowed to have, as well as the participant's inherent non-physical attributes and the attributes of the environment itself. The function $h(\cdot)$ that maps attributes into emotional satisfaction is complex. As a result,

designing such things as online game worlds is a difficult task. It seems likely, however, that aggregate well-being can only rise as more of these worlds appear.

IV. Inequality and Avatar Society

In the preceding discussion, we have assumed that the characteristics of a given virtual reality space are fixed, from the standpoint of the choosing agent. In truth, the attributes of the world are open to influence by agents, via three routes.

First, in some worlds, the software explicitly allows agents to use their avatars in ways that permanently change the environment. In *Norrath*, the world of the game *EverQuest*, certain avatars may cast spells that stop the rain. Other worlds (*Shadowbane*) allow avatars to build and destroy cities. Thus, the physical attributes of the world may be directly changed by avatars.

Second, the experiences that a given world provides depend on the society of other agents who participate in the world. Agents, through their avatars, will occupy certain social roles, and these roles may provide better or worse experiences depending on the relative abundance or scarcity of other avatars occupying the same role. As an example, in *Brittania*, the world of the game *Ultima Online*, avatars can make shirts and sell them. Once upon a time, a certain agent decided to collect and hoard 10,000 shirts. As the price of shirts rose, those who happened to be accomplished shirtmakers with some stock on hand did quite well. If that agent then opted to give away the hoard, the shirtmakers would have all faced unemployment. As in Earth economy and society, the return to doing any given thing depends on the collective valuation of that thing, as well as the relative scarcity of other people doing it. In this sense, agents may change the

social, economic, and cultural attributes of a virtual reality environment, in the same way that they influence the social, economic, and cultural attributes of Earth.

Third, the attributes of a world can be changed by the coding authority (the game owners, for example), and agents have a number of methods by which to induce the changes they desire. For example, the world of EverQuest allows agents to inhabit a certain class of avatar known as druids. Once upon a time, the life of a druid at a certain point was limited by the existing rules of the game. The druids mounted an effective quasi-political campaign to have the rules changed, and after months of petitioning by email and on discussion boards, the coding authority for EverQuest (Verant Interactive), gave the druids what they wanted.⁹ Verant is part of a competitive company, Sony, so it can be assumed that the decision was motivated to a large extent by the pursuit of profit. The incentive of Sony to pursue profit allowed the druid community to pressure the company for change. Thus, the company's decision was the outcome of a meta-game between the participants of the world and the owners. The owners control the coding authority, and the coding authority can make the world into absolutely anything that the mind can imagine; therefore, anything that the agents can do to influence the owners, in some kind of game outside the game, can be used to change the very nature of the world.

In these cases, the characteristics of a virtual world can be changed by individual or collective decisions of the agents, both inside and outside the virtual world itself. Modeling the choice of avatars and worlds, then, must recognize that world characteristics \underline{z} are endogenous, that there are N agents playing the meta-game (N is very large), and that the action space covers actions within all worlds, both virtual and Earthly. Because of the evident complexity of this problem, it would be fruitless to

attempt a general theory. To do so would be equivalent to building a general N-player theory of the game of human social, cultural, economic, and political change. Yet, to even note this equivalence is itself striking, because it suggests that the problems of designing and then living in shared virtual reality environments are really no different from the ancient human problem of designing and then living in our Earth environment. Earth must have a society, an economy, and a politics, and so must a place like Norrath. The emergence of these difficult choice problems in cyberspace has nothing to do with the fact that human agents are interacting via avatars in virtual reality; it has everything to do with the fact that they are human agents, interacting. The attributes of the worlds of humans have always been endogenous.

That having been said, there are some critical differences between world management on Earth and world management in cyberspace, and some of them can be explored theoretically. The most important difference involves the ability of the coding authority to define the attributes of avatars. On Earth, our attributes are determined at birth, for the most part.¹⁰ Some of these attributes are permanent, and inhere in the choosing agent forever. Others, however, exist only in the context of Earth, and can be dramatically different in other worlds. No government on Earth can ensure that all its citizens will have equal height; no government on Earth can ensure that all citizens will have equal intelligence. Similarly, no coding authority can ensure equal intelligence of all agents who drive avatars in its virtual world. However, any coding authority can easily guarantee that all avatars, driven by whomever, will have the same height. Thus, by allowing more control over the distribution of some of the most basic attributes that

people have, involving even the very nature of their physical beings, virtual worlds permit the shaping of human society at a deeper level than ever before.

Indeed, at first glance, the opportunities for constructive social engineering might seem endless. Granted that no one can change the attributes \underline{x} that are inherent in agents, is it possible to design a world such that the attributes \underline{y} of all avatars are equal? This would remove from the calculus of social inequality all effects derived from such bodily characteristics as height, weight, age, gender, clothing, hair color (and amount!), beauty, and sexual attractiveness. Removing the inequality that derives from these things might dramatically reduce the overall inequality in well-being. Anecdotal evidence from existing virtual worlds suggest that many who go there are, in fact, on the lower end of Earth's distribution of well-being, often because of some purely physical factor.¹¹ As it turns out, removing these inequalities is trivially easy to do; very early avatar spaces such as Habitat (Morningstar and Farmer, 1991) were populated by avatars who were exactly the same in all respects other than appearance. The emergence of these physically-equal virtual worlds would seem to offer opportunities to not only enhance average well-being, but also to reduce its variance.

Developments in this market suggest, however, that very few people are interested in living in worlds where all avatars are equal. Put bluntly, life in such worlds seems to be boring. At any time, participation in highly-equal worlds like AlphaWorld or Habitat is dwarfed, by several orders of magnitude, by participation in highly-unequal worlds like Norrath.¹² This suggests that humans' attitudes toward the construction of social inequality are quite complex. To suggest that there is a general, monotonic, aversion to inequality (as in Atkinson, 1970) is very much at odds with observed

behavior here. Rather, there seems to be a desire for some inequality, or a certain kind of inequality, but not too much.

In sum, the theoretical issues are as follows. Profit-seeking companies build virtual worlds, in the hopes of obtaining revenues from either participation fees or advertising. It is generally profitable to have a large number of participants.¹³ They therefore endow these worlds with attributes that, they hope, will attract large numbers of people. The people, on the other hand, seek to spend time in worlds that provide them with emotional satisfaction. The data indicate that there is much more demand for worlds with a great deal of avatar inequality, such as Norrath, than for worlds with no avatar inequality, such as AlphaWorld. Therefore, the worlds that agents seek to create for themselves seem to involve social diversity and inequality as a core element. Unequal societies are apparently a typical of the meta-game among the agents.

V. Inequality in a Model of Avatar Production

To better understand the reasons why worlds of avatar inequality have appeared, it will be useful to describe and model the way inequality is generated in Norrath and other game-based virtual realities. The social worlds of these games are built around three common principles that apparently contribute strongly to their popularity.

The first principle is division of labor: Agents seem to desire avatars with unique abilities, by which they can provide individualized contributions to avatar society. The second principle is equality of opportunity: Agents seem to enjoy a rags-to-riches storyline, in which everyone starts out very weak and very poor, but then has the opportunity to advance through the application of time and skill to gameplay. The third

principle is inequality of outcomes based on merit only: Agents seem to prefer game mechanisms that grant advantages of wealth and power only to avatars who have performed more meritorious actions (where “merit” is admittedly hard to define – working long hours at the game, being socially or politically skillful, etc.). Together, these three principles attempt to provide diversity, equality, and meritocracy, and this seems to be the most desired kind of society.

An example makes the issues more concrete. Suppose a new player in a game must choose to be either a warrior or a wizard. The two avatar types are very different in their bodily characteristics. Warriors are strong but cannot control magic, while wizards are brilliant spellcasters but physically weak. Warriors begin with rusty weapons and armor; wizards begin with a simple spell that produces a fireball the size of a sparkler. By fighting monsters and performing quests (which is doing work, effectively), warriors can obtain funds to buy better armor, and wizards can obtain funds to buy better spells. With more investment in the avatar, the avatar becomes richer and more powerful. The player may choose either the warrior’s or the wizard’s path, and the ideal game would have the two paths be different in form, but potentially equal in outcome. If, after investing perhaps 500 hours of labor, both warriors and wizards can have about the same amount of wealth and power, then the game has provided a diversity of social roles, with an equality of opportunities. And, if it so happens that a given warrior invests 500 hours while her friend the wizard invests only 100 hours, then it is acceptable that the warrior has more wealth and power than the wizard. She worked for it. Games that meet all three principles in this way tend to be very popular; we can surmise that these societies are considered by many people to be ideal.

Constructing societies that meet these principles well, however, is extremely difficult. It is perhaps easiest to provide diversity, and indeed most new virtual worlds seem to eschew the Habitat model of equal abilities. Instead, worlds are populated with many different avatar races, and within each race, an avatar can have a number of different professions or occupations. Moreover, after birth, an avatar can often work on various skills (baking, tailoring, archery) to further specialize. Gameplay is then built on teams, where the ideal is to put together groups of well-matched, diverse avatars, to accomplish goals that they all have in common. However, as diversity increases, it seems to become more difficult to ensure equality of opportunity and merit-based advancement. Within the game, it becomes difficult to ensure that all races and professions have equal access to income and new powers. Outside the game, it becomes difficult to prevent people from using Earth resources to advance themselves. Designing a world that comes reasonably close to meeting all three principles thus seems to involve introducing enough diversity to provide individuality, while at the same time managing the experiences and powers of the players so that no one has an unfair advantage.

The design problems are actually much simpler than they seem, however, when viewed in the long run. This can be illustrated in a simple general equilibrium model of avatar production. Suppose there are two types of avatars available in a world, type 1 and type 2. Let N_1 and N_2 , respectively, indicate the population of avatars of each type. By “population” is meant the total amount of time, within a fixed time period, that agents devote to avatars of each type. The analogy here is to aggregate labor supply as a productive factor in a macroeconomic model of national production.¹⁴

Total production of valued goods and services in the virtual economy is given by $y = f(N_1, N_2)$. The production function is continuous, twice-differentiable, with $f_j > 0$, and $f_{jj} < 0$, $j = 1, 2$, and $f_{12} = f_{21} > 0$. The positive cross-derivatives of $f(\cdot)$ indicate that the marginal product of wizards will be higher if there are more warriors around. This reflects the idea that avatars do better when working in diverse teams, which, in turn, is the standard way of applying the principle of division of labor in the existing virtual worlds.

If we assume a competitive factor market, each hour of avatar labor will be compensated according to its marginal product. Let $g_j = \partial f / \partial N_j$ be the goods and services obtained per hour of input by avatar type j . From the standpoint of the choosing agent, g_j is the wage earned for each hour spent driving a type j avatar. Let there be I agents, indexed i , with $i = 1, \dots, I$. Each agent contributes an amount n_{ij} of each type of avatar labor, $j = 1, 2$. Hence $N_1 = \sum_i n_{i1}$, and $N_2 = \sum_i n_{i2}$. Each agent's share of total production is $c_i = g_1 n_{i1} + g_2 n_{i2}$.

We assume that each agent has already solved a higher-level maximization problem in order to determine the total amount of time to devote to all avatars in this particular world, which we will consider fixed. For simplicity, suppose that each agent has chosen to devote the same amount of time to this world; denote it T . This allocation will be spent on the two avatar types, so $T = n_{i1} + n_{i2}$. Agent utility derives from her payoffs from production in the game world, c , according to a continuously twice-differentiable function $U(c)$, with $U' > 0$ and $U'' < 0$. The agent's choice problem is

$$\text{Max}_{n_{i1}, n_{i2}} U(g_1 n_{i1} + g_2 n_{i2}) \text{ subject to } T = n_{i1} + n_{i2}$$

In the Earth economy, different input types are generally supplied under differing cost conditions, so that the opportunity cost of supplying factors generally differs by factor. In virtual economies, however, an agent can supply either type 1 or type 2 avatar hours and the opportunity cost of doing so – the sacrifice of her time – is the same regardless of the factor supplied. Therefore, if $g_1 > g_2$, each agent will devote all her time to avatar type 1. If $g_1 < g_2$, each agent will devote her time to avatar type 2. If $g_1 = g_2$, each agent will be indifferent between the two types; any time allocation provides the same level of well-being.

On the production side of the economy, equilibrium in the factor market requires that the marginal rate of transformation in production is equal to the ratio of returns:

$$\frac{\partial f(N_1, N_2) / \partial N_2}{\partial f(N_1, N_2) / \partial N_1} = \frac{g_2}{g_1} \quad (1)$$

A general equilibrium requires that the wages g_1 and g_2 adjust so that condition (1) holds, and, simultaneously, each agent makes an individually optimal allocation of time to the two avatar types.

Suppose that the economy is initially in a state where $g_1 = g_2$, and condition (1) holds. In that case, the aggregate avatar populations are at levels N_1^* and N_2^* with the property $\partial f(N_1^*, N_2^*) / \partial N_1 = \partial f(N_1^*, N_2^*) / \partial N_2$. Such a state is feasible. Because the wages are equal, each agent is indifferent between the two avatar types. Any allocation of labor between the two types is therefore consistent with consumer rationality. We can assume that agents simply mirror the aggregate allocation. More precisely, if $n_0 = N_1^* / (N_1^* + N_2^*)$, then each agent can allocate $n_{i1} = n_0 T$ hours to avatar type 1.¹⁵

Now suppose that something perturbs the economy away from this state. Perhaps the coding authority changes the production function $f(.,.)$ so that, at the current

allocation of avatar types, $g_1 > g_2$. All agents now have the incentive to abandon their current mixed allocation and play avatar 1 exclusively. As a thought experiment (so as to avoid modeling the dynamics explicitly), imagine that agents act on these new incentives over time.¹⁶ As agents react to the wage difference, they begin to devote more time to type 1 avatars. As they do so, the marginal productivity of type 1 avatars falls, and the marginal productivity of type 2 avatars rises. This is because of the team-based nature of the production function; the influx of new warriors makes warriors comparatively less valuable, while making wizards comparatively more valuable. As productivity changes, so do wages: g_1 falls and g_2 rises. This dynamic of population shifts and wage changes will continue as long as g_1 exceeds g_2 . Thus, as long as g_1 exceeds g_2 , g_1 will fall and g_2 will rise. Eventually, the two wage rates will become equated, but now in a population that has a larger share of type 1 avatars.

By similar argument, whenever g_2 exceeds g_1 , g_2 will fall and g_1 will rise, until a new equilibrium is established with a higher population of type 2 avatars. The only stable point – the only long-run equilibrium – is found when $g_1 = g_2$. Thus, equality of wage rates and marginal productivity across avatar types is an inherent characteristic of long run equilibrium in these economies.¹⁷

Define “equality of opportunity” as the state of affairs that exists when $g_1 = g_2$. Then we can also say that an avatar economy with diversity can also be one with equality of opportunity, so long as a) the production function which regulates the wages of the avatars rewards specialization and the division of labor, and b) agents are free to shift hours of time from one avatar type to another. The idea is that differences in the productivity of avatar types cannot be sustained in the long run. If, for some reason, one

avatar type were to become more powerful or productive than another, agents would flood the world with new avatars of that type. Whatever advantage those avatars have at first would be eroded as they compete with one another for increasingly scarce positions in the economy. At the same time, the other, disadvantaged avatar types would find an increasing number of the advantaged type to work with, thus raising their own ability to be productive in the economy. A new equilibrium would eventually be established, with more of the ‘advantaged’ avatar type than before, but with wage equality restored. Equality of opportunity is something inherent in a team-based avatar production economy; the economy seeks equal wages on its own, through the operation of its internal long run dynamics.

The implication is that managing diversity, equality, and meritocracy in virtual worlds may be much simpler than it seems. One can simply create diversity, and then structure the production of value in the economy on the basis of teams. That is, make every avatar type valuable, in some way, to the others. Then the marginal productivity of every avatar type will rise when that type becomes relatively scarce, and also when other avatar types become relatively abundant. Population shifting will automatically equate the wages of each kind of avatar. Once the returns to avatar types are equal, the system seems largely meritocratic: those who have more power in the virtual world have simply put more time into their avatars. In essence, virtual worlds provide human beings with the opportunity, to slip costlessly from one physical identity to another. As a result, inequalities among physical entities will be eroded by basic competitive pressures, and differences in outcomes will come to be determined more by mental effort than by the bodily endowments represented by the avatars.

This is a happy outcome, of course, but there is a reason why it does not work on Earth: on Earth, we cannot easily shift from a body with one set of characteristics to some other body with a different set.

Several exceptions to this generally optimistic normative conclusion need to be discussed. First, these statements are about the adjustment of avatar production systems in the *long run*. In virtual worlds where agents spend hundreds of hours developing avatar capital, the long run can seem very long indeed. It is fairly common for coding authorities to suddenly change rules that have a dramatic impact on the marginal productivity of different avatars. Often, these rules changes are justified in the name of exactly the kind of equality of opportunity issues that actually seem to be taken care of automatically by the system. For those who already have extensive prior investments in avatars of a given type, these abrupt changes can quickly and dramatically reduce the value of these investments. While the system will eventually restore the balance among the classes, it will not necessarily compensate specific individuals for their losses. Agents threatened by possible rules changes typically try to generate a great deal of political heat to protect themselves. By the same token, it is fairly common for agents to assume, initially, that a certain class of avatar is inherently quite powerful in a game, only to see its net advantage erode as more and more agents enter it. Agents in this situation also generate political heat, lobbying for rules changes that would keep their old advantages intact. All of these political pressures are quite similar to the political pressures we see on Earth: it may be that we are all equal in the long run, but politics operates in the short run too. The production system of a virtual economy is therefore often subject to sudden shocks that redistribute well-being. What this model suggests is that game designers

would be well-advised to resist short-run political pressures for immediate changes, and, instead, allow equality of opportunity to restore itself gradually over time.

Second, it may be the case that avatars may differ so much in their attributes that the production function cannot provide equal payoffs under any allocation of the population. Perhaps *everyone* finds avatars of type X quite boring to play. The model of this section is a simplification of the model of section IV, in that it assumes utility is derived entirely from the production of the avatars. While that production may be thought of in material terms (loot), it can also be thought of as units of emotional satisfaction (“funs” or “hedons”). The model, however, assumes that utility is a function of all the attributes of the avatars and the worlds, and it may well be that certain avatars are so much more, or less, fun to play that no amount of shifting of player populations will ensure equal payoffs in a team-based model of hedonic production. In effect, this is really just to caution that corner solutions are always possible.

A third caveat involves what is perhaps the most difficult of the three principles to ensure, that being merit-based advancement. In the model, agents increased their production and consumption only by devoting time to their avatars. On the surface, such inequalities in outcomes seem entirely meritocratic and one might predict that worlds based on these kinds of productive systems would be perceived without envy by most participants. In practice, however, the outcome inequalities that exist in virtual worlds do raise considerable concerns among agents. One reason is that some agents resent the inherent abilities and endowments of others. In the context of this model, one can imagine that the time endowment T might be different among different agents, and this might lead to differences in outcomes that seemed unfair.¹⁸ Another source of resentment

is that some agents apply resources to the investment problem that are said to be improper. For example, some agents use their Earth incomes to purchase powerful avatars and special equipment on online auction sites like Ebay.¹⁹ To many, this is a violation of the spirit of the game, which requires that agents act as though they live only in the virtual world. It is important to maintain the fantasy atmosphere of a world, and doing so requires that resources earned outside the game world not be applied there. The issue is largely one of values, but the point is that managing agents' attitudes about the source of investments in avatar capital, represents an outgoing challenge to the coding authorities. Of the three principles, merit-based advancement seems hardest to ensure.

Overall, however, this simple model seems to suggest that the emergence of virtual worlds may be destined to reduce the level of human social inequality. The existence of multiple worlds in which to live allows people to sort amongst the worlds; the existence of multiple body-types in a given world allows people to sort amongst the body-types. In the long run, this sorting should produce equality of opportunities for well-being across avatar types and worlds, and therefore across individuals. The equality of opportunity will apparently not be brought about by forcing an equality of attributes; instead, agents seem to prefer worlds in which avatars have a very broad spectrum of differentiated talents and abilities. Rather, it is the long-run dynamics of population shifts among differentiated avatars that equates the rewards available to each avatar type. Equality of available rewards, however, will not necessarily result in equality of outcomes, as agents may apply their resources with greater or lesser intensity. There will continue to be intense debates about social questions, both because short-run inequalities may persist for some time, and because certain kinds of investments and endowments

may seem unfair if used in the context of a particular world. In the end, social and political debates about inequality may be just as heated in virtual worlds as they are on Earth, but their subject matter will derive less and less from inequities in our bodies – simply because the citizenry can easily migrate from one body to another.

VI. Identity, Reputation, and Social Norms

In a universe that offers hundreds of virtual worlds, the body becomes a vessel of choice, and the thinking part of humanity – the self – will find it convenient to slip into and out of avatars as economic, social, and political circumstances dictate. At first glance, this may seem to pose a distinct challenge for economic theory. Economics is built on the idea that the self is a unitary, rational actor. Psychologist Sherry Turkle (Turkle, 1995), however, has argued that the ability of the agent to represent herself as a different person in different online communities, without anyone being able to trace one identity to another, effectively creates multiple ways of knowing, which can be thought of as multiple selves. Perhaps this is just a semantic issue; what are referred to as “multiple selves” in such an argument are not the same as the “unitary, rational, choosing self” of economic theory. Rather, the “multiple selves” only have their existence because a unitary higher-order actor, choosing rationally, decided to create and then inhabit them. This higher order actor is the “self,” or the “agent,” in the models developed in the preceding sections. In fact, one could argue that the emergence of anonymity on the internet changes nothing fundamental about the nature of human behavior. For centuries, improvements in technology have allowed the self to act in different ways in different communities, without anyone being the wiser. The internet only intensifies this ability.

What is changing is not the Self, which remains unitary, but the ease with which the Self can manipulate its manifestation in various physical spaces.

Perhaps the emergence of avatars will reveal behaviors that seem contradictory under current theories about the nature of tastes. But again, this is nothing new. People have lived double lives for centuries. The sense in which these lives are “double” is entirely a social construct, however. In 1968, it was widely believed that no one over the age of 30 could possibly enjoy rock and roll. If one closely followed a number of old folks around, however, one would certainly have discovered some who listened to Lawrence Welk on Thursday but went to a Beatles concert on Friday. They did not tell their friends about the Beatles concert only because, from the standpoint of society’s judgments, such tastes would have been judged “irrational” and “inconsistent.” But it is the individual mind that determines what tastes cohere. From the standpoint of economic theory, no contradiction is indicated when someone appears in Norrath as both a young man and an old woman. Indeed, if variety is the spice of life, economic theory *predicts* that the unitary actor will choose a number of different physical manifestations by which to appear. The development of avatars, and the shifting of the self between them, has no real consequence for the applicability of rational choice theories.

These developments do have implications for the communities that humans form, however. Rational choice theories of social effects (Bernheim, 1994; Binmore and Samuelson, 1994; Young, 1993) emphasize the importance of information for the maintenance of social norms. Norms can only be enforced if it is possible to impose some kind of punishment on the violators. In a virtual community, the real self behind the avatar is generally hidden. As a result, any punishments the community may dictate can

only be imposed on the avatar, not the self, and the self is free to simply exit the avatar and escape unscathed. For example, unfair play in a pay-for-play game like EverQuest is punished by banning a user from the game; the user's account is closed, and all her avatars effectively executed, permanently. However, nothing can stop the banned user from opening a new account, with a different credit card, and starting new avatars. Thus, it would seem that there is nothing to prevent anyone from violating any and all social norms, without consequence. The prospects for stable community in such an environment seem bleak. The instability of online communities has long been a subject of study for sociologists (Smith and Kollock, 1998).

However, the economic literature also suggests (Eshel, Samuelson, and Shaked, 1998) that agents will sort themselves into different groups depending on how interested they are in living in a community regulated by social norms. Such sorting is evident in existing virtual worlds, in that worlds with built-in systems for maintaining user reputations seem vastly more popular than worlds where reputations cannot be known. For example, AlphaWorld, as mentioned above, endows all avatars with the same abilities at all times. Thus, anyone who violates a social norm in AlphaWorld, if banned, can immediately create a new avatar, with a different name, but with all of the same abilities and powers as before. The community can have no effect on behavior. Contrast this with a game like EverQuest. In EverQuest, a user's ability to be a grief to others depends on her level of abilities, and abilities can only be gained by devoting hours to an avatar, in team-based operations with other avatars. Advancement in the game therefore requires that a user become known for good play, so as to be accepted into teams. A player who violates norms will not be able to advance very far, simply on grounds of

reputation; in fact, there is little or nothing that a user can do in EverQuest without the help of others. A user may consider obtaining a new reputation by simply starting a new avatar, but new avatars are so weak and poor that they can do very little harm to anyone. The combination of team production and level-based advancement seems sufficient to support very strong social norms in the game, and such norms do seem to be present.²⁰

The broader implication is that the diversity of avatar attributes is useful not only as an end in itself, but also as a means to encourage agents to develop and maintain good reputations for their avatars. Specialized attributes allow systems of joint production, which can have the side effect of inducing conformity to social norms.

At the same time, no norm system in cyberspace can be truly oppressive; anyone who wishes to live under a certain set of norms, or under no norms at all, is free to inhabit avatars in the appropriate worlds. However, it should be noted that the same systems that encourage norm formation will also tend to slow down the shifting of populations that leads to equality of opportunity. Equality of opportunity emerges because dissatisfied agents can develop capital in different kinds of avatars; it will emerge more quickly if the switch from one powerful avatar to another can be accomplished rapidly. Yet a system that encourages reputations will also require that it is *not* easy to develop capital rapidly in another avatar. If it were easy to do so, anyone could ruin their reputation with one avatar, destroy it, and then simply reappear with another avatar of similar powers. The credibility of social norm enforcement depends on the degree to which an agent has a vested interest in the fate of the avatar. If agents are deeply invested in their avatars, and are hesitant to start new ones, it will take more time for the leveling process of population

shifting to occur. There is a trade-off between equal opportunity and social order; agents will choose worlds based on their relative tastes for both.

VII. Conclusion: The Avatar and Public Policy

This paper has developed explored the nature of avatars from the standpoint of simple economic theory. An avatar is the manifestation of a rational, choosing self in a particular physical environment, virtual or “real.” The demand for an avatar is a demand to spend time “being” or “driving” or “inhabiting” it, and this demand is derived from the attributes of the avatar, the world the avatar exists in, and the self. The demand also depends on the choices of other agents, as these choices can change the nature of the world the avatar inhabits. In particular, the relatively scarcity or abundance of different types of avatars can have an effect on the value of a single avatar’s time. In long run general equilibrium, though, the shifting of selves among avatars will tend to eliminate any disparities in the net return to an hour of time spent driving any given avatar. Avatars will tend to yield similar returns for the same reason that competing investments offer similar returns: activities that generate profits attract entry; those that generate losses induce exit. While entry and exit will encourage equitable returns among different forms of avatar capital, it may also cause social degradation, in that shifting selves are harder to bring into conformity with local norms of acceptable behavior. There is some tension between individual interest and the common good.

The paper also developed a series of normative points, suggesting that avatar-based living seems likely to have a net positive effect on well-being. Before the advent of the avatar, there was only one world to live in, Earth, and only one avatar to inhabit there,

the Earthly body. The recent emergence of virtual worlds besides Earth has vastly expanded the range of choices regarding one's own physical being and the space which it inhabits. For those who are heavily burdened by their Earth bodies – because of physical challenges, or because of the often-brutal social stigma that is applied to physical appearance – the opportunity to inhabit a different body must generate an increase in well-being. For anyone, the opportunity to have different bodies at different times allows easy experimentation with a variety of social roles – warrior, dancer, mentor, prostitute – and if variety is the spice of life, an expansion of these opportunities also must raise well-being. In general, the development of shared virtual reality environments has provided a significant expansion in the number of different lives a person may lead. The expansion of choice sets in this fashion can only raise aggregate well-being.

There is also reason to believe that the cross-section variance of well-being will fall. If people are given greater opportunities to shift among different bodies, the impact of disparities in body type on the disparity of well-being will decline. No one can be trapped in a body that, for whatever reason, provides a comparatively low rate of social, economic, and political rewards to its owner. Instead, bodies that are comparatively unattractive will become comparatively rarely seen, since few agents will choose to drive them. If these bodies provide anything at all to the community, that service will be eventually be in comparatively short supply, and will therefore be comparatively well-rewarded. By the same argument, bodies that provide many services to their owners will be abundant, and therefore poorly-rewarded. In the long run, an agent should be able to enjoy about the same degree of economic, social, and political rewards, regardless of the characteristics of the body she inhabits.

This is not to argue that the life of an avatar will be perfect or blissful in any sense, of course. The distribution of assets, resources, and talents on Earth will have an influence on the ability of different agents to produce desired outcomes through their avatars. Mental abilities, financial endowments, and dumb luck on Earth will continue to affect human well-being. As a result, even the best coding of a virtual world will not produce a blissful, stress-free paradise. Social and political pressures will simply revolve around issues other than those posed by the body. The argument here is not that such pressures and problems will disappear, but that they will manifest themselves on a social plateau that provides more well-being, and a fairer distribution of it, than is now the case on Earth. Future research is needed to determine how specific actions by the coding authority and others affect the behavior and well-being of agents.

From the broadest perspective, the argument above relies on a number of unstated assumptions about public policy. The most critical unstated assumption has been that agents are more or less free to choose their avatars and the worlds those avatars inhabit. These freedoms allow for competitive markets for avatar time and avatar production, which should raise well-being and remove some of the inequities of Earth living. Implicitly, these assumed freedoms also give agents a role in defining the kind of worlds they experience. Agents are assumed to have the power to build social communities with unique norms of behavior, and to protect those communities from external interference.

If these assumptions are violated, the emergence of virtual worlds may not have the happy consequences foreseen here. The film *The Matrix* supposes that “the machines” have become the sole providers of sensory inputs to human brains, and have thereby trapped and enslaved all human bodies. Far-fetched as this may seem, there is no question

that those who write the code for virtual worlds can obtain power over those who become immersed there. Some refer to games like EverQuest as “addictions,” because life there is so compelling and enjoyable that the agent finds it difficult to leave. As the technology of sensory immersion expands, it is possible that an ever larger share of the population will spend an ever larger fraction of its time receiving sensory input from a shared virtual reality environment. In that state, they may well be vulnerable to exploitation.

The principle protection that contemporary humans have against exploitation derives from their rights as citizens. Because the avatar-mediated life is fairly unique, however, it makes sense to be specific about what rights need to be protected in order to ensure that avatar-mediated living will be a good thing. First and foremost, it needs to be recognized that the rights in question adhere to *agents*, not avatars.²¹ The agent is where the conscious mind resides, and it is the mind, and not the body, that is the ultimate possessor of human dignity. If the beneficial effects of virtual living are to be assured, the rights of agents that should be protected include:

- *Freedom of movement and association.* Agents should be free to leave and enter virtual worlds as they wish. Agents should be free to form their worlds into unique communities. They have the right to protect themselves and their communities from harassment and abuse.
- *Freedom of information, both incoming and outgoing.* Agents may receive a very large share of their information from a virtual world. They have a right to expect that that information is not fraudulent or harmful. They have a right to send and receive information from and to multiple sources, including those outside the virtual world in which they spend most of their time.

- *Freedom of speech, assembly, and political participation.* Decisions that affect the well-being of agents with avatars in a virtual world, are inherently political decisions. The agents therefore have a right to participate in the formulation of policies that affect them. This extends to the writing of code. The coding authority is subservient to the will of the people.²²

One might argue that these rights are already guaranteed to most citizens of advanced countries, and their maintenance in the context of virtual worlds will therefore end up being a fairly trivial issue. A counter-argument would point to current practice in the majority of virtual worlds, where all three of these rights are routinely violated. Indeed, in most virtual worlds, the user is confronted with an End User Licensing Agreement that abrogates these rights as pre-condition of participation. In EverQuest, for example, the users must agree that all property they create in the game belongs to Sony (the game's operator). They also agree that they will not use any third-party program in connection with the game, meaning that they cannot, for example, add a program that might allow them to surf the web in a separate window while playing. And they agree to conform to acceptable behavior; the company reserves the right to ban any user for improper behavior, at its discretion. Similarly, the end-user agreement for Earth and Beyond (by Westwood/Electronic Arts) prohibits users from "disparaging the Product, EA, Westwood or any of their representatives or employees." There is no guarantee that a company's notions of acceptable disparagement will be legitimate in the same way that a democratic society's notions would be. Moreover, users who are in the virtual world live under heavy restrictions on the flow of information. Such conditions are more than likely

to prevent – within the world, at least - the development of political movements that might legitimize, or oppose, the policies of the coding authority.

In the end, the users of the current generation of large-scale virtual worlds have only one source of political influence: Exit. So far, when combined with the profit motives of the companies, exit seems to have encouraged policies that keep the virtual worlds compelling, immersive, and fun, despite the fact that the agents involved have given away a number of fairly critical entitlements. At the same time, conflicts between users and companies are almost constant, and usually revolve around the issues covered by the three rights listed above; users may complain, but they really have no choice but to accept the company's judgment in any dispute, or leave the world.²³ Only time will tell whether a system grounded in the competitive market for games will be the best way to govern these new societies.

If not, it remains unclear by what authority the rights of agents can be guaranteed. Agents can come to a virtual world from anywhere on Earth, and it is not clear whether Earth governments, Earth companies, or the community of agents itself should be the ultimate forum of appeal.

References

- Atkinson, A.B. (1970), "On the Measurement of Inequality," *Journal of Economic Theory*, 2, 244-263.
- Becker, David (2002), "Game Exchange Dispute Goes to Court," at news.com: <http://news.com.com/2100-1040-832347.html?tag=rn>.
- Becker, Gary (1996), *Accounting for Tastes*, Cambridge, Harvard University Press.
- Bernheim, B. Douglas (1994), "A Theory of Conformity," *Journal of Political Economy* **102**(5), October, pp. 841-877.
- Binmore, Ken and Larry Samuelson (1994), "An Economist's Perspective on the Evolution of Norms," *Journal of Institutional and Theoretical Economics*, v150 n1, March, pp. 45-63.
- Castronova, Edward (2001), "Virtual Worlds: A First-Hand Account of Market and Society on the Cyberian Frontier," CESifo Working Paper No. 618, December.
- _____ (2002), "On Virtual Economies," CESifo Working Paper No. 752, July.
- Eshel, Ilan, Larry Samuelson, and Avner Shaked (1998), "Altruists, Egoists, and Hooligans in a Local Interaction Model," *American Economic Review*, 88(1), March, 157-179.
- Morningstar, Chip and F. Randall Farmer (1991), *The lessons of Lucasfilm's Habitat*, in Michael Benedkit, editor, *Cyberspace: First Steps*, Cambridge: MIT Press.
- Patrizio, Andy (2002), "Did Game Play Role in Suicide?", April 3, at wired.com: <http://www.wired.com/news/games/0,2101,51490,00.html>
- Pham, Alex (2002), "Software Comes Back in Play as Game Industry Seeks New Winners," *Los Angeles Times*, May 20.
- Pontifical Council for Social Communications (2002), *Ethics in Internet*, at the Vatican: http://www.vatican.va/roman_curia/pontifical_councils/pccs/documents/rc_pc_pc_cs_doc_20020228_ethics-internet_en.html
- Smith, Marc A. and Peter Kollock (eds.) (1998), *Communities in Cyberspace*, New York: Routledge.
- Turkle, Sherry (1995) *Life on the Screen: Identity in the Age of the Internet*, New York: Simon and Schuster.
- Young, H. Peyton (1993), "The Evolution of Conventions," *Econometrica* **61**(1), January.

Endnotes

¹ According to the Los Angeles Times (Pham, 2002), game industry revenues were \$9.35 billion in 2001, of which \$3 billion was for hardware. Total Hollywood box office revenues in 2001 were \$8.4 billion.

² ‘Avatar’ is a Hindu term for an incarnation of Vishnu. It was first applied to representations of people in computer graphics by the makers of Habitat (Morningstar and Farmer, 1991).

³ In truth, with current technology, a person participating in an internet-based shared virtual reality environment can be receiving sensory input from two sources at once. To see things in the virtual reality space, one looks at the computer screen. The screen displays an image of the VR space as it would appear to the eyes of the avatar. At the same time, peripheral vision and momentary sideward glances also deliver images from Earth reality to the mind. Similarly, the mind receives sound from both cyberspace and Earth; it receives tactile sensation exclusively from the Earth. Thus, given current technology, “inhabiting an avatar” actually involves nothing more than receiving significant sensory input from a shared virtual reality space. As technology advances, the degree to which signals from cyberspace can dominate total sensory reception will increase.

⁴ In November 2001, Shawn Wooley apparently committed suicide over events in Norrath, the game world of EverQuest (Patrizio, 2002). Players of Lineage have been known to hunt down and kill one another. These isolated cases illustrate the extent to which virtual reality can occupy emotional space in the mind. The far more common anecdote, from my experience, is the woman who reports losing a relationship over gameplay. “I could not get him to turn his head away from the computer” is a common statement.

⁵ More concretely, a computer program registers a sensory input to the avatar in cyberspace, and that input is delivered via internet, and then interface devices (monitors, speakers, force-feedback gloves, etc.) to the sensory receptors of the mind.

⁶ For the time being, assume that all world attributes are physical attributes, and hence fixed. Later, this assumption will be relaxed in order to explore the social attributes of worlds, which of course can be altered by the collective choices of the agents.

⁷ This assumption is already violated at times. A minority of players of role-playing games jack in to the game using two computers at once, and drive one avatar on each. To keep things simple, I will continue to assume that each mind can drive at most one avatar at a given time.

⁸ Again, to be specific: the sensory input of a “tall ogre” in the virtual world would differ from that of a “short human” in Earth, in that the virtual world’s computer program would detect the sensory images of the world as though they were being received by a twelve-foot-tall being. It records the position and height of an ogre’s eyes and then renders the view that eyes in that position and height would perceive. It then sends those images directly to the agent’s mind. In this way, the agent perceives the virtual world as if she were twelve feet tall. Note that this may or may not mean that she experiences the virtual world as a person who is taller than the person she inhabits on Earth – the virtual world may consist entirely of fifteen foot tall ogres. If she happens to like being taller than others, in this case her utility function would put as much or less weight on her ogre avatar as on her Earth avatar, as both leave her short. More on this below.

⁹ I refer to the druid campaign to obtain the Complete Healing spell ability, which was successfully completed when the ability was granted in September, 2002.

¹⁰ True, with an investment of time and money, a person can change some attributes of her Earth avatar (Becker, 1996). Yet the will and vision and resources to make the change are themselves attributes. In its most precise expression, my argument here is that there are certain attributes that a person on Earth can change only at great cost, if at all, but that a coding authority in cyberspace can change at a trivial cost.

¹¹ This is the analog to a theoretical idea presented in another paper (Castronova, 2002), which suggests that people with low wages on Earth would find virtual worlds comparatively more attractive. Similarly, those who, on Earth, are very overweight or cannot walk, may well find life in a virtual body much more enjoyable. If a virtual world were to offer an equality of body weight and walking speed, those who feel stigmatized on those grounds in the context of Earth society would find that world attractive.

¹² An unequal world like Norrath typically has more than 60,000 people in it at any one time. Visit AlphaWorld (probably the most famous and popular of the non-game avatar spaces) and you will see perhaps eight people, or as many as 20 or 30.

¹³ However, there are congestion effects. The profit-maximizing population of a virtual world is not infinite. See Castronova, 2002.

¹⁴ In existing game worlds, the most accurate concept would not be avatar population but “avatar capital.” This would be the aggregate of avatar hours weighted by the skill levels of the avatars. The analogy is to human capital, which could be measured, roughly, by the aggregate labor supply weighted by the hours of education of each worker. Avatar capital depends on skill levels because avatars of higher levels are able to produce more goods and services.

¹⁵ Note that equality of marginal productivity does not imply equality of populations. This state of affairs would be consistent with having the world dominated by warriors, with only a few wizards, or vice versa.

¹⁶ It would be conceptually straightforward, but space-consuming, to express this problem in terms of avatar capital accumulation over time, with a dynamic general equilibrium framework. That having been said, a dynamic general equilibrium is much closer to reality: in practice, agents build avatar capital slowly and by many strategies, from developing the skills of a given avatar, to buying powerful avatars outright, for cash, on online auction sites like Ebay.

¹⁷ This result depends, of course, on the team-based nature of the production function. If a given avatar type is of no use to others, then there is no guarantee that population shifts will make its wage equal to that of any other avatar.

¹⁸ This kind of complaint is usually leveled by more mature players, who have families, jobs, and children, against younger players, who often seem to have nothing to do besides play the game. In the long run, the younger players typically have much more powerful avatars, and this is deemed unfair.

¹⁹ Typically, it is the more mature agents who spend real money on their avatars. For those who “have a life,” as they put it, cash is the abundant resource. For those who don’t “have a life,” the abundant resource is time. It is economically rational that the old and the young make use of the relatively abundant resource to make their investments in avatars. Indeed, the dollar-based market for avatar capital is the direct result of the comparative advantage that high school and college students have in the time-consuming activity of building avatar capital within virtual worlds. They trade with older people, whose comparative advantage is in the earning of dollar-denominated incomes from the avatar capital they have developed on Earth.

²⁰ This statement is based on anecdotal evidence only. As a player, I have often experienced the Norrath equivalent of “jumping on the hand grenade,” actions which only make sense if the individual is concerned about long run reputation effects.

²¹ There has already been a declaration of the Rights of Avatars (by Raph Koster, see <http://www.legendmud.org/raph/gaming/index.html>), which stresses that the issues involve the well-being of the players behind the avatars. Similarly, the Catholic Church holds that all social communication, including the internet, should be “...by persons to persons for the integral development of persons” (Pontifical Council for Social Communications, 2002).

²² One might also suggest that property rights are fundamental, but I would argue that they are already covered in the rights listed. There are already several real-world cases involving property rights and virtual worlds (Becker, 2002). They arise because, in many games, players eagerly seek out opportunities to sell the goods they have acquired on online auction sites, such as Ebay. While this practice is consistent with the right to property, it violates the right to community: Most of these games are based on a rags-to-riches storyline, and if someone can buy his way out of virtual poverty using Earth riches, those who cannot do so lose the value and enjoyment of their gaming experience. The community of agents has the right to protect their world and its atmosphere. Perhaps mindful of the long-run negative effect on profits of a decline in game atmosphere, the companies that own the games often declare these activities to be “against the rules” and banish any guilty players from the game world. Thus, at the moment, the conflict is being played out as if it is between the Ebay-selling players versus the profit-seeking company. In the final analysis, however, this is a case of conflict between the rights of individuals and the rights of the community.

²³ On August 30, 2002, the gaming site Arcadian del Sol reported the following story. Two users of the Ultima Online virtual world Britannia were discussing an event in their home town, and one said “anyone who commits rape should be castrated.” A passer-by was offended by the statement, captured a log of the conversation, and sent it to customer support. The owning company, Electronic Arts, has given its workers a clear and unambiguous rule: Those who use the word ‘rape’ in a sentence will have their accounts banned. And so they were. The agents behind the banned avatars had probably devoted hundreds of hours to the creation of valuable assets in Britannia. By banning the accounts, the company effectively destroyed the fruits of these labors, without compensation or right of appeal. See

<http://www.arcadiandelosol.com/article.php?sid=141&mode=thread&order=0>. Such stories of capricious adjudication are commonplace.

CESifo Working Paper Series

(for full list see www.cesifo.de)

- 798 Mariam Camarero, Javier Ordóñez, and Cecilio Tamarit, The Euro-Dollar Exchange Rate: Is it Fundamental?, October 2002
- 799 Misa Tanaka, How Do Bank Capital and Capital Adequacy Regulation Affect the Monetary Transmission Mechanism?, October 2002
- 800 Jörg Baten and Andrea Wagner, Autarchy, Market Disintegration, and Health: The Mortality and Nutritional Crisis in Nazi Germany, 1933-1937, October 2002
- 801 Saku Aura, Uncommitted Couples: Some Efficiency and Policy Implications of Marital Bargaining, October 2002
- 802 Wolfram F. Richter, Delaying Integration of Immigrant Labor for the Purpose of Taxation, October 2002
- 803 Gil S. Epstein and Shmuel Nitzan, The Politics of Randomness, October 2002
- 804 John Hassler and José V. Rodríguez Mora, Should UI Benefits Really Fall over Time?, October 2002
- 805 Friedrich Breyer and Stefan Felder, The Dead-anyway Effect Revis(it)ed, October 2002
- 806 Assar Lindbeck and Solveig Wikström, E-exchange and the Boundary between Households and Organizations, November 2002
- 807 Dieter Bös, Contests Among Bureaucrats, November 2002
- 808 Steven Brakman, Harry Garretsen, and Marc Schramm, The Strategic Bombing of German Cities during World War II and its Impact on City Growth, November 2002
- 809 Florian Englmaier and Achim Wambach, Contracts and Inequity Aversion, November 2002
- 810 Sarbajit Sengupta, Delegating Recruitment under Asymmetric Information, December 2002
- 811 Rajshri Jayaraman, On the Partial Public Provision of a Private Good, December 2002
- 812 Stéphanie Stolz, Banking Supervision in Integrated Financial Markets: Implications for the EU, December 2002
- 813 Christian Keuschnigg, Taxation of a Venture Capitalist with a Portfolio of Firms, December 2002

- 814 Inés Macho-Stadler and David Pérez-Castrillo, Settlement in Tax Evasion Prosecution, December 2002
- 815 Rainer Niemann and Dirk Simons, Costs, Benefits, and Tax-induced Distortions of Stock Option Plans, December 2002
- 816 Jan-Egbert Sturm and Barry Williams, Deregulation, Entry of Foreign Banks and Bank Efficiency in Australia, December 2002
- 817 V. Anton Muscatelli, Patrizio Tirelli, and Carmine Trecroci, Monetary and Fiscal Policy Interactions over the Cycle: Some Empirical Evidence, December 2002
- 818 Claude Hillinger, A General Theory of Price and Quantity Aggregation and Welfare Measurement, December 2002
- 819 Erkki Koskela and Ronnie Schöb, Optimal Capital Taxation in Economies with Unionised and Competitive Labour Markets, December 2002
- 820 Sheilagh Ogilvie, Guilds, Efficiency, and Social Capital: Evidence from German Proto-Industry, December 2002
- 821 Hans Gersbach and Verena Liessem, Financing Democracy, December 2002
- 822 Costas Hadjiyiannis, Panos Hatzipanayotou, and Michael S. Michael, Optimal Tax Policies with Private-Public Clean-Up, Cross-Border Pollution and Capital Mobility, December 2002
- 823 François Ortalo-Magné and Sven Rady, Homeownership: Low Household Mobility, Volatile Housing Prices, High Income Dispersion, December 2002
- 824 Syed M. Ahsan and Panagiotis Tsigaris, Measuring the Social Discount Rate under Uncertainty: A Methodology and Application, December 2002
- 825 Kai A. Konrad, Altruism and Envy in Contests: An Evolutionarily Stable Symbiosis, December 2002
- 826 Robert S. Chirinko and Huntley Schaller, A Revealed Preference Approach to Understanding Corporate Governance Problems: Evidence from Canada, December 2002
- 827 Geir B. Asheim, Green National Accounting for Welfare and Sustainability: A Taxonomy of Assumptions and Results, December 2002
- 828 Andrea Gebauer, Chang Woon Nam, and Rüdiger Parsche, Lessons of the 1999 Abolition of Intra-EU Duty Free Sales for Eastern European EU Candidates, December 2002
- 829 Giacomo Corneo, Work and Television, December 2002
- 830 Vivek H. Dehejia and Yiagadeesen Samy, Trade and Labour Standards – Theory, New Empirical Evidence, and Policy Implications, December 2002

- 831 Geir B. Asheim and Wolfgang Buchholz, A General Approach to Welfare Measurement through National Income Accounting, December 2002
- 832 Aaron Tornell and Frank Westermann, The Credit Channel in Middle Income Countries, January 2003
- 833 Gebhard Flaig, Time Series Properties of the German Monthly Production Index, January 2003
- 834 Campbell Leith and Jim Malley, Estimated Open Economy New Keynesian Phillips Curves for the G7, January 2003
- 835 Burkhard Heer and Bernd Süßmuth, Inflation and Wealth Distribution, January 2003
- 836 Erkki Koskela and Leopold von Thadden, Optimal Factor Taxation under Wage Bargaining – A Dynamic Perspective, January 2003
- 837 Carola Grün and Stephan Klasen, Growth, Income Distribution, and Well-Being: Comparisons across Space and Time, January 2003
- 838 Robert S. Chirinko and Ulf von Kalckreuth, On the German Monetary Transmission Mechanism: Interest Rate and Credit Channels for Investment Spending, January 2003
- 839 Sascha O. Becker, Andrea Ichino, and Giovanni Peri, How Large is the “Brain Drain” from Italy?”, January 2003
- 840 Albert Berry and John Serieux, All About the Giants: Probing the Influences on Growth and Income Inequality at the End of the 20th Century, January 2003
- 841 Robert Fenge and Martin Werding, Ageing and the Tax Implied in Public Pension Schemes: Simulations for Selected OECD Countries, January 2003
- 842 Robert Fenge and Martin Werding, Ageing and Fiscal Imbalances Across Generations: Concepts of Measurement, January 2003
- 843 Giovanni Andrea Cornia, The Impact of Liberalisation and Globalisation on Income Inequality in Developing and Transitional Economies, January 2003
- 844 Peter Fredriksson and Per Johansson, Program Evaluation and Random Program Starts, January 2003
- 845 Bernd Hayo and Matthias Wrede, Fiscal Equalisation: Principles and an Application to the European Union, January 2003
- 846 Syed M. Ahsan and Jaideep Oberoi, Inequality, Well-being and Institutions in Latin America and the Caribbean, January 2003
- 847 Chang Woon Nam and Doina Maria Radulescu, The Role of Tax Depreciation for Investment Decisions: A Comparison of European Transition Countries, January 2003

- 848 V. Bhaskar and Steinar Holden, Wage Differentiation via Subsidised General Training, January 2003
- 849 Paloma Lopez-Garcia, Labour Market Performance and Start-up Costs: OECD Evidence, January 2003
- 850 Christian Keuschnigg and Soren Bo Nielsen, Public Policy for Start-up Entrepreneurship with Venture Capital and Bank Finance, January 2003
- 851 Yin-Wong Cheung, Menzie D. Chinn, and Eiji Fujii, China, Hong Kong, and Taiwan: A Quantitative Assessment of Real and Financial Integration, January 2003
- 852 Gregory D. Hess, The Economic Welfare Cost of Conflict: An Empirical Assessment, February 2003
- 853 Douglas J. Cumming and Jeffrey G. MacIntosh, Comparative Venture Capital Governance. Private versus Labour Sponsored Venture Capital Funds, February 2003
- 854 Eckhard Janeba and John Douglas Wilson, Decentralization and International Tax Competition, February 2003
- 855 Tapio Palokangas, Capital Accumulation and Employment Cycles in a Model of Creative Destruction, February 2003
- 856 Brendan Walsh, When Unemployment Disappears: Ireland in the 1990s, February 2003
- 857 Luis H. R. Alvarez and Erkki Koskela, A General Approach to the Stochastic Rotation Problem with Amenity Valuation, February 2003
- 858 Christian Schultz, Strategic Campaigns and Redistributive Politics, February 2003
- 859 Ernst Fehr and Joseph Henrich, Is Strong Reciprocity a Maladaptation? On the Evolutionary Foundations of Human Altruism, February 2003
- 860 Haizhou Huang, Dalia Marin, and Chenggang Xu, Financial Crisis, Economic Recovery and Banking Development in Former Soviet Union Economies, February 2003
- 861 Pedro Cardoso and Bernard M.S. van Praag, How Sustainable Are Old-age Pensions in a Shrinking Population with Endogenous Labour Supply?, February 2003
- 862 Volker Meier, Efficient Transfer of Aging Provisions in Private Health Insurance, February 2003
- 863 Edward Castronova, Theory of the Avatar, February 2003