

ON VIRTUAL ECONOMIES

EDWARD CASTRONOVA

CESIFO WORKING PAPER NO. 752

CATEGORY 9: INDUSTRIAL ORGANISATION

JULY 2002

An electronic version of the paper may be downloaded

- from the SSRN website: http://ssrn.com/abstract_id=338500
- from the CESifo website: www.CESifo.de

ON VIRTUAL ECONOMIES

Abstract

Several million people currently have accounts in massively multi-player online games, places in cyberspace that are effectively large-scale shared virtual reality environments. The population of these virtual worlds has grown rapidly since their inception in 1996; significantly, each world also seems to grow its own economy, with production, assets, and trade with Earth economies. This paper explores two questions about these developments. First, will these economies grow in importance? Second, if they do grow, how will that affect real-world economies and governments? To shed light on the first question, the paper presents a brief history of these games along with a simple choice model of the demand for game time. The history suggests that the desire to live in a game world is deep-rooted and driven by game technology. The model reveals a certain puzzle about puzzles and games: in the demand for these kinds of interactive entertainment goods, people reveal that they are willing to pay money to be constrained. Still, the nature of games as a produced good suggests that technological advances, and heavy competition, will drive the future development of virtual worlds. If virtual worlds do become a large part of the daily life of humans, their development may have an impact on the macroeconomies of Earth. It will also raise certain constitutional issues, since it is not clear, today, exactly who has jurisdiction over these new economies.

JEL Classification: L86.

Keywords: information and internet services, computer software.

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

“I play for my kid. He's a cool dude and has some health problems that keep him out of school from time to time and he needs something to focus on. Don't get me wrong, this kid was born being 80 years old, incredibly mature, creative and intelligent enough to manage A's even when attending school only 3 out of 5 days for most of the year. Anyways, we've always been ‘simpatico’, both being creative and analytic at the same time. He loves to share his on-line adventures. What started 3 years ago with Tribes (yes, I was a clan mom) evolved into Tribes RPG (yes, we ran a mod on our own server) and segued into DAoC [Dark Age of Camelot]. After hearing him talk about it, and scrolling through forums, I decided this looked intriguing. It would be a good way to ‘buddy up’ and go have some fun with him. Anyways, I watched over his shoulder one day, and said, ok, that's it, I'm getting a version myself. Started that night. D&D/RPG stuff is fairly easy to get one's head around. We teamed up, and proceeded to level quickly...My daughter of course, had to join the fun. Soon the whole family was talking about strategy and ‘the funny thing that happened today while I was leveling...’” – Gelenii, posted 6/8/02 at daoc.catacombs.com.

I. Introduction

At this writing, there are several million people around the world who have access to a MMORPG account.¹ A MMORPG (massively multi-player online role playing game), more simply a *virtual world*, is an internet-based game that can be accessed by large number of players at the same time. Players choose a physical self, an avatar, and then spend their time running about in the game world, chatting with others, undertaking various tasks, purchasing, producing, and consuming goods, and generally leading a more or less full, rich, and detailed life there. Many of these players spend no more time in virtual worlds than they do in ordinary hobbies. Many others, however, approach virtual worlds as an alternative reality, devoting a substantial fraction of their time to them. According to a survey in Summer 2001, about one third of the adult players of EverQuest spent more time in a typical week in the virtual world than in paid employment

¹ A rough accounting as of June 2002, based on media reports and company claims: Lineage, 4 million; EverQuest, 400,000; Ultima Online 200,000; Camelot, 200,000; Anarchy Online 200,000; Asheron's Call 150,000. This does not include any number of chat-based games and some smaller titles that exist. On the

(Castronova, 2001a). Since that time, at least ten major new titles in development have been announced, including several by corporate powerhouses such as Microsoft, Vivendi, and Sony. As this market expands, it seems entirely possible that living a part of one's life in cyberspace may eventually become a common practice.

Such a development would be worth some attention, because life in cyberspace seems to be different in important ways from life on Earth. This is especially true of economic life. True, at first glance there are many similarities between Earth economies and their virtual counterparts. In an earlier paper (Castronova, 2001a), I described the economy of Norrath (the virtual world of the game EverQuest) as if it were a normal Earth economy, complete with statistics covering such activities as production, labor supply, income, inflation, foreign trade, and currency exchange. There is evidence that the economies of these virtual worlds generate a surprisingly high level of per capita production, and that people who 'live' there (a substantial fraction view themselves as citizens) have accumulated significant stocks of real and financial wealth. All of this suggests that there is something very normal and mundane about cyberspace economies; people live there, work there, consume there, and accumulate wealth there, just as they do on Earth.²

other hand, many people have accounts in several games. No one knows how many people are actually active in game worlds, but the number appears to be substantial.

² Some of the concepts and ideas in the paper may seem unusual to readers who are not familiar with the basic aspects of the online gaming world; my earlier paper would be a decent introduction (Castronova, 2001a). The one finding of that paper that deserves added emphasis here is that players in virtual worlds have developed durable assets of considerable economic value. On a per capita annual basis, the production of new value (in terms of both durable equipment and avatar capital, i.e. the skill set of the avatars) is on par with that of a moderately advanced industrial country. The pseudo-GDP of Norrath amounted to as much as \$2,000 per capita annually. There is also a substantial dollar-based trade in goods, as well as a liquid market in currencies. In short, even a brief and limited empirical assessment revealed evidence of real economic activity, both within online virtual worlds, and between virtual worlds and earth. The present paper takes this economic activity as a given, and tries to explore some of its broader implications.

However, further thought suggests that virtual economies may be anything *but* normal. As an example, consider a simple policy question: Should governments attempt to control prices? Most economists would say ‘no,’ since the costs of doing so outweigh the benefits. Moreover, the costs often end up being borne by the people the policy is supposed to help. These perverse effects happen because any effort to control prices creates either excess supply or excess demand, which in turn generates all kinds of social costs. Surplus goods must be bought up and destroyed, or shortage goods must be allocated by a mechanism that usually turns out to be both unfair and costly. But what if it cost the government nothing to buy up a surplus of goods and destroy it? And what if the government could simply produce whatever quantities are demanded, at no cost to itself? If those two acts were possible, then a policy of government price control would be feasible. And in cyberspace, the coding authority does indeed have the power to create and destroy any amount of any good, at virtually zero cost. Therefore, as a de facto government, the coding authority can indeed control prices. And therefore, price controls may actually be good policy in cyberspace, even though they most certainly are not good policy on Earth.

The example suggests the possibility that virtual economies may be very different from Earth economies, in certain well-defined ways. As economic and social activity gradually migrates from Earth to cyberspace, these differences may begin to have an impact on the lives of large parts of the population. Details about the functioning of virtual economies may, in time, become important public issues. Even today, small changes in the code of a game can generate intense controversy among the players. If these little firestorms are a portent of things to come, it would be useful, even now, to

analyze some of the unique features of virtual economies and ask how these features may eventually influence economic and public policy questions.

As an initial approach to these issues, consider the following two kinds of questions:

1. *The future of games*: Will multi-player online games become an important part of the social life of humans? What does the market for games look like? What sort of market structure can we anticipate in the future?
2. *The impact of games*: How would a large emigration of work and play time to these virtual worlds affect the economy of the real world?

The paper attempts to address these questions, beginning, in Section II, with a brief history of games. This is not so much a general history of gaming as an effort to gauge the cultural importance of avatar games – games in which the player embodies a game piece in the same way that avatars tend to become a person's bodily self in the virtual world. Section III then develops a simple rational choice model for determining the demand for game time. Given the unique features of gaming as a consumer good, Section IV assesses possible market structures in the games market - will one game eventually dominate the world? Section V considers the macroeconomic implications of large-scale expansion in the gaming phenomenon, especially for GDP and the tax base. Section VI lays out some of the policy issues that widespread gaming will raise. Section VII concludes with a list of simple teachings in economics that are held to be always true on Earth, but that seem to be less than entirely true, or at least open to doubt, in virtual worlds. These topics represent avenues of future research.

II. A Brief History of Avatar Games

Participation in virtual worlds is, at the moment, participation in a market for games. At an abstract level, what are the special features of such a market? How do games differ from other goods in the economy? The first step would be to define the object we are after, but defining “game” is difficult. The Oxford English Dictionary lists 17 senses of the word, from an “amusement, diversion, or pastime” to “a proceeding, scheme, intrigue, undertaking, followed up like a game.” Perhaps it is easier to say that our culture identifies certain practices specifically as games, and we are most interested in the class of games that have now become virtual worlds. That is, while recognizing that both solitaire and football are games, they are not games that have evolved to the point that the people who play them view them as essentially an alternative reality. I will refer to games with some sort of alternative reality as ‘avatar games.’

In an avatar game, the players use a single game piece to represent themselves in the make-believe play environment. *Avatar* is a Sanskrit word that identifies earthly manifestations of the god Vishnu. When Vishnu incarnated himself as a tortoise, he was named Kurma; when he appeared as the fish, he was named Matsya.³ The idea is that the Earthly being – the avatar tortoise or fish - was the embodiment of a higher being. In 1985, F. Randall Farmer and Chip Morningstar developed Habitat, the first multi-user domain with a visual 2D interface, and they chose *avatar* as the term for the cartoons that users would drive around the virtual world. Thus I, Castranova, would have appeared in the Habitat multi-user environment as “Bird,” a parrot-headed puppet. All of the contemporary virtual worlds use this convention; when you enter the world, you are

³ Information on Vishnu and other Hindu teachings from the Columbia Encyclopedia, 6e.

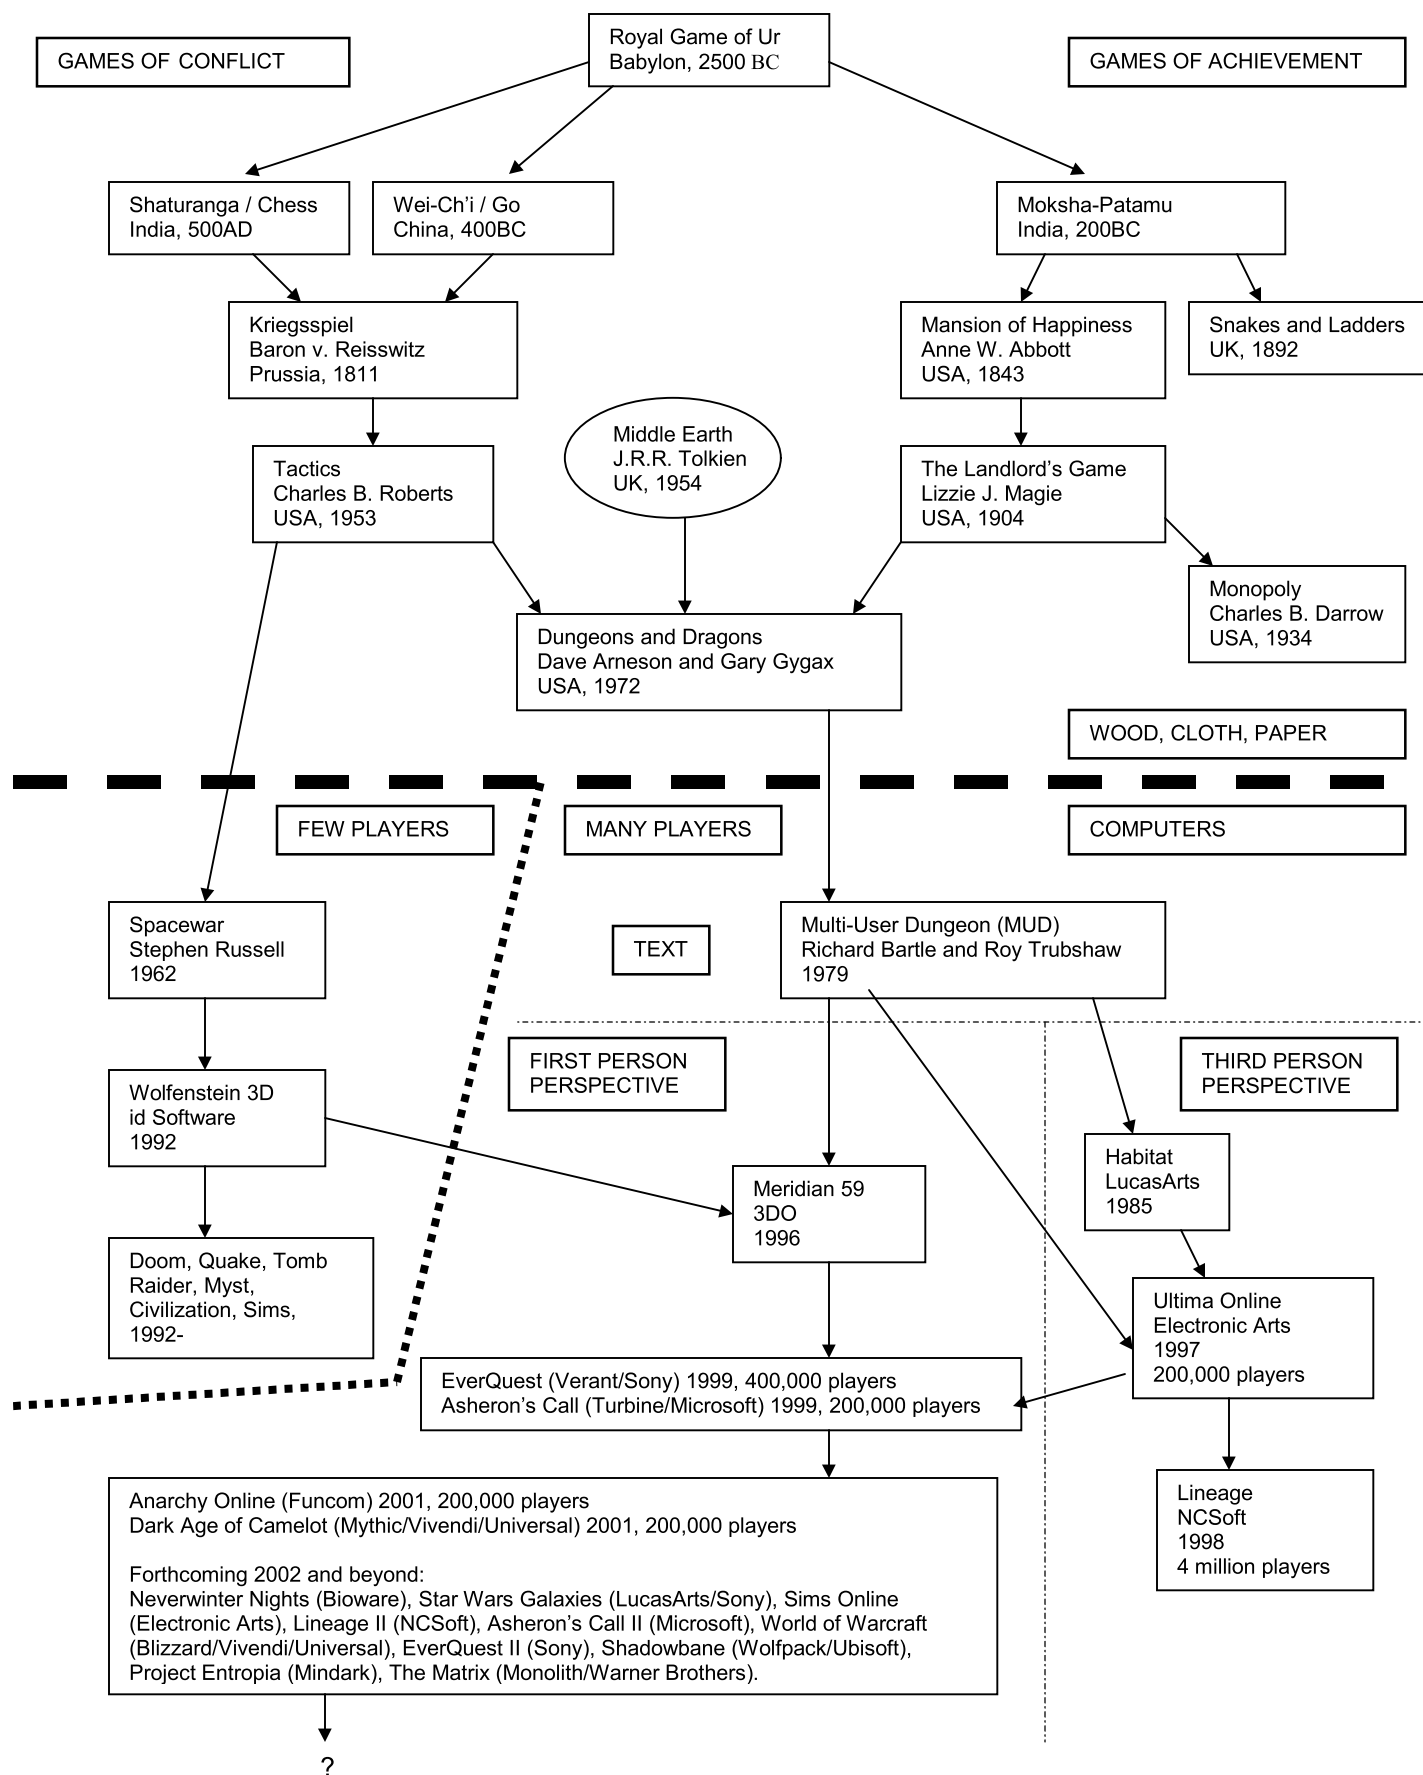
driving a 3D representation of a being in that space; that being is your avatar, the embodiment of you in that physical environment.

History suggests that avatar games such as these are by no means a transient cultural phenomenon. On the contrary, they appear to be an extremely ancient aspect of culture. Moreover, human interest in avatar games seems to have been driven by technology: as technology has enhanced the immersive experience of the games, the games have become more popular.

These connections are illustrated in Figure 1, which provides a short timeline of game development from ancient games using wood, paper, and cloth, to the most recent computer-based role-playing games.⁴ The first game that we know of was the Royal Game of Ur (c. 2500BC), and it seems to have been an avatar game; it involved racing a game piece - the first game avatar - around a board. Some 2000 years after the Royal Game of Ur, we have the first games identifiably related to games we know of today. Games simulating warfare seem to have been the first to emerge, the most famous being chess and go. Advances in technology allowed these conflict and war games to become more and more realistic (Kriegsspiel, 1811), so that by the mid-1960s it was possible to play a number of games offering fairly accurate simulations of actual battles (e.g. Tactics, 1954). By the mid-1970s, there were any number of games which allowed a person to immerse himself for many days in the fantasy of being Napoleon or Hitler.

⁴ Most of this information is general knowledge; some of the specifics, especially dates and names, come from James Masters' online guide to traditional games (<http://www.tradgames.org.uk>). The lines of influence are not direct or causal; they represent my own beliefs about what elements in the culture seemed to allow the next step in gaming development.

Figure 1. Development of Avatar Games



Meanwhile, a parallel development saw the racing aspect become a metaphor for personal achievement. Moksha-Patamu (c. 200 BC) is apparently the first game in which a player embodies a single playing piece and then attempts to develop or build that playing piece into something better. I would argue that it is the ur-ancestor of all avatar games. The game is known today as Chutes and Ladders, and was intended to teach Hindu children how to attain moksha (heaven, bliss, and release from the eternal cycle of reincarnation) despite being tempted by various sins (patamu). Doing good deeds allowed players to advance on ladders, doing bad deeds dropped them down chutes (snakes in the Hindu original). Moksha-Patamu was imported to the UK in 1892 as Snakes and Ladders, but its game mechanic was applied to morality games much earlier in the Victorian era, notably Anne W. Abbot's *Mansion of Happiness* (1843). Later, Lizzie J. Magee applied the avatar game mechanic to the acquisition of wealth in the *Landlord's Game* (1904). Charles B. Darrow took the same ideas and made *Monopoly* in 1934, a game that soon became standard in household game cupboards all over the world. As with conflict games, by the 1960s these individual-advancement games had developed to the point that a person could immerse himself in the fantasy of being a landlord, tycoon, or other kind of mogul for hours or days at a time.

At this point, two entrepreneurs in the simulation wargaming subculture, Dave Arneson and Gary Gygax, hit upon the idea of making a simulation at the 1:1 level, so that each player commanded exactly one soldier. And rather than simply endow their soldier with a standard set of equipment, they opted to allow a soldier to build up his powers and skills through good play, as in moksha-patamu. In order to broaden the skills

a soldier could attain, they then chose a fantasy setting, drawing on concepts whose ultimate source was J.R.R. Tolkien's Middle Earth. The result was a set of gaming rules that was released as "Dungeons and Dragons" in 1972. As is well known, Dungeons and Dragons offered a very immersive experience; an urban legend at the University of Wisconsin (known to the author) holds that several old D&D players have maintained their alternative reality in the campus tunnels for decades.

By this point, computers have begun to enter the picture. The first video game was a wargame, Spacewar, invented by Stephen Russell in 1962. There then emerged a whole genre of popular computer-based conflict games for one player or a small team. In 1992, id software released Wolfenstein 3D, a game with 3D graphics and a first-person perspective. For the first time, the player actually could see the world through the eyes of an avatar.

Over the same time period, computer games with many players emerged, initially as text-based multi-user domains on internet-linked terminals. Richard Bartle and Roy Trubshaw developed the first of these, basing the game on Dungeons and Dragons. In 1985, F. Randall Farmer and Chip Morningstar made Habitat, a multi-user domain using 2D avatars roaming in 2D spaces, in a third-person perspective. What was most interesting about these MUDs was the fact that they did not go away when the player stopped playing. The game environment and the objects in it persisted in time. Players could choose to visit the place whenever they wished. The only real difference between these places and our world was the lack of a first-person visual perspective: you could live in that place, but you could not perceive it through your own eyes. At best, you could see yourself there; you could not actually *be* there.

In 1996, 3DO combined the MUD multi-player idea with a Wolfenstein-like 3D graphics engine, and released Meridian 59, the first virtual world per se. The game had the immersive feel of a 3D first-person action game, but the game world persisted in the same way that MUD worlds persisted. In 1997, Electronic Arts combined the MUD concept with Habitat's 2D world, and released Ultima Online, again a virtual world that did not go away when the player left the game. While Meridian 59 did not catch on with the playing public to any great extent, Ultima Online did, and still has at least 200,000 active subscribers. Ultima Online then spawned Lineage, a 2D multiplayer game that has become the most popular virtual world, with some 4 million subscribers.

In 1999, Sony's EverQuest and Microsoft's Asheron's call managed to build very large and successful virtual worlds using the first-person 3D perspective of Meridian 59. Both games are still a strong presence in the market, with hundreds of thousands of active subscribers. First-person 3D seems to be the future of the genre; the two releases of 2001, Dark Age of Camelot and Anarchy Online, both were first-person 3D. Both games have been successful, again with hundreds of thousands of subscribers within a few months of opening. As for the future, the final box in Figure 1 lists several virtual world titles that are known to be in development by well-funded companies. The market in massively multiplayer first-person 3D games seems poised to grow.⁵

⁵ There are two strange coincidences in this history. First: Chip Morningstar, an inventor of Habitat, coined the term "avatar" to refer to the 2D cartoon figures that would be used to represent users in the Habitat environment. He did so apparently unaware that Habitat's ultimate ancestor is moksha-patamu, the game that illustrates the Hindu virtues and vices. As legend has it, those virtues and vices first came to humanity's attention as such when they were revealed to us by Krishna, who was himself an avatar of Vishnu. Thus, the concept of "avatar" in current games seems to have appeared without any prior consciousness of the historical association of these games with a real Hindu avatar. Second: Gyax and Anderson, the inventors of Dungeons and Dragons, made heavy use of the concepts of character classes and levels, and these concepts continue to play a very important role in contemporary games. Players choose a class (a type of person to be) and then try to raise their level (their competence or human capital) by devoting time to the development of skills that are specialized to that class. This advancement structure

Overall, the historical perspective delivers a fairly strong message: human participation in avatar gaming seems to be driven strongly by technological evolution. Our desire to live a fantasy life has probably been constant, although one could argue that industrialization, the mechanization of existence, the growth of the work ethic, and the resulting destruction of family life may have made the desire for make-believe living more intense. But the real contribution of technology has been in the means by which this fantasy life becomes possible. Wooden game pieces do not provide a particularly immersive experience, and it takes a great deal of mental effort to maintain the fantasy that one really is bobbing up and down between Moksha-Patamu's heaven and earth. But now that computers have begun to make near-immersion experiences possible (if you have ever flinched while playing Quake, you will know what I mean), the amount of time and resources devoted to producing these experiences seems to have been rising.⁶ This suggests that as the technology of gaming advances, we can expect gaming to become an increasingly important part of daily life.

Does any of this make games worthy of deeper exploration? When people enter a virtual world, perhaps they are "just playing a game," so why should it matter. On the other hand, we can also say the same thing when people "play" the stock market, a game that certainly does matter. Nonetheless, to some, there is a difference between interactions that we label, literally, as games (checkers) and ones that we label as games only in a figurative sense (the stock market; elections). And then it is claimed that

is similar to certain aspects of Hindu religion, in which a person is allowed to choose one of several "ways," with each way entailing a certain path of advancement to enlightenment. Arneson and Gygax, however, got their ideas from the specialization of roles in western military formations, not from Hindu teaching. It is odd that several core concepts of Hinduism seem to have appeared unintentionally and independently, in two places, in the world of 21st century gaming.

behavior in the literal games is more important than behavior in the figurative games. To this way of thinking, there is a difference between real life and game life, and if something is only a game, it isn't "real." So it does not matter.

To others of a more contemporary mode of thinking, however, the line between real and unreal is blurred. When post-modernists say, "life is a game," they are not just invoking a metaphor, they mean it. Life really is a game, and games really are life.⁷ If one is important, so is the other. Games matter.

The fact that games matter to postmodern cultural theorists does not mean they deserve any specific treatment by economists. To address this issue, the next section asks what an economic theory of gaming behavior might look like, and what the consequences of gaming may be for economic well being.

III. An Economic Theory of Games: The Puzzle of Puzzles

To develop a theory of the games market, it might be natural to start with the field of game theory. However, game theory, as a research program, is mostly interested in improving methods of general strategic analysis – its objective is not to analyze the

⁶ We cannot know for sure whether game participation is higher on a per capita basis today than in 2500BC. However, it does seem to be the case that children today spend more time playing video games than children of an earlier era spent playing games like Monopoly.

⁷ Umberto Eco's *Foucault's Pendulum* (Eco, 1989) makes a fairly simple point: If everyone believes in some conspiracy theory, it is possible for society to make it true, for all intents and purposes. Widespread belief in the theory imposes social constraints that eradicate all behavior inconsistent with the theory. Simple self-interested incentives could ensure, for example, that no one would ever doubt a theory stating that Martians came to earth in 1800 and decreed that Americans would drive on the right side of the road while Britons would drive on the left. Anyone attempting to live at odds with the theory would be killed right off. The broader point is that the truth and reality are slippery concepts. Jorge Luis Borges' "Lottery in Babylon" (Borges, 1962) uses Game as a metaphor for this social process: the lotto begins as a game but it becomes social life in toto. My point in bringing this up is not to argue that Truth and Reality do, or do not, exist distinctly from Fantasy and Game. My point is that the fluidity of these concepts in actual human minds may (indeed already do) allow millions of people to slip fairly easily from one world to another. Friday's accountant is Saturday's wizard. This behavior has observable economic, social, and political consequences, for people on either side of these lines.

markets for cultural objects identified by the word “game.” Game theory is interested in fascinating games like the Prisoner’s Dilemma, but you will not find a home version of that game in toy stores (“Now with repeated N-player action!”). And a search of several literatures in the social sciences suggests that neither game theorists nor anyone else has devoted much time to the things that we call games in the real world, even ones as simple as checkers.⁸

Absent any specific prior theoretical treatment in economics, perhaps the most intuitive approach might be to think of the market for games as a market for simple, durable entertainment goods. From the modernist way of thinking, this would be the right choice: there is nothing in a game but entertainment, which people enjoy and pay for; hence the demand for games is like the demand for, say, books.

⁸ In June 2002, I searched the Econlit database (which covers for articles in economics, political science, public policy, and elsewhere) for the following terms: checkers, chess, go (in titles only; ‘go’ in keywords brings too many hits), landlord’s game/darrow (‘monopoly’ returns too many hits), jeopardy, let’s make a deal, backgammon, cribbage, tetris, contract bridge (‘bridge’ returns too many hits), yahtzee, tomb raider, sim city, euchre, pac-man, trivial pursuit, myst, craps, poker, blackjack, slots/slot machines, and horse racing. The results: zero hits for all games except: chess (2), jeopardy (1), let’s make a deal (1), contract bridge (1), poker (2), slots (2), and horse racing (14). The seminal work on chess (Simon and Schaeffer, 1992) argues persuasively that games as complex as chess are not usefully studied from a von Neumann – Morgenstern game theory perspective: there are so many nodes in the tree that it is impossible to make an exhaustive assessment of the terminal value of any given move. A massively multiplayer online roleplaying game like EverQuest is infinitely more complex than chess; it is not even clear what ‘victory’ means. Beyond Simon and Schaeffer (1992), few of the other papers address games *qua* games. The game-show papers are fairly tongue-in-cheek analyses of the incentives posed by certain aspects of these games (Metrick, 1995; Page, 1998). The bridge and poker papers are in the vein of ‘Here is a phenomenon that is a lot like bridge/poker,’ not really analyses of the games as actually played (Shubik, 1999; Nash and Shapley, 1997; Engwall, 1994; Mazalov, Panova, and Piskuric, 1999). The slots papers are about gambling revenues (Nichols, 1998). Only the horse racing papers amount to a legitimate literature on the game in question, as actually played in real life. Yet horse racing is actually a spectator sport, a subject for which we do have a well-developed literature and journals (e.g. the *Journal of Sports Economics*). On the whole, then, it is safe to say that social scientists simply have never explored the games that people really play. (There does seem to be a fairly large literature in the AI community.) The oversight is especially glaring in the case of games like Sim City and its spin-offs. I suspect that many of my colleagues in the social sciences have played Sim City, enjoying the job of Big City Mayor for an evening. Like them, it never occurred to me to ask whether the fact that millions of other people also seemed to like role-playing this job was worthy of further investigation. Nor did it occur to me that this behavior, in itself, might reveal something fundamental about the economic and social behavior of people. The opportunity to make these investigations is only growing, however; those mobs of pseudo-mayors will soon become part of actual virtual cities in the Sims Online.

The post-modernists would counter that at the level of massively multiplayer gaming, the metaphor of games-as-books breaks down fairly completely. Gaming remains an entertainment good, but it immerses the player so thoroughly in the virtual society and economy that events in the virtual world have an emotional impact on people no different from the impact of Earth events. Events in the virtual world can have an influence that extends well beyond the borders of the virtual world; relationships, incomes, and even lives on earth may be affected. Thus, a post-modern way of thinking would require a theory that is unique to games, in recognition of the importance that real-world people place on events that happen in the game.

Economists are apparently forced to be post-modernists on this question. The economic theory of value seems to require it. In economics, the value of objects does not depend on their characteristics or their components, but rather on their contribution to the well-being of the people who use them. Value is subjective, wholly created in the minds of people. If people in free markets determine that a shiny crystal called ‘diamond’ is worth \$100,000, economists basically accept the reality of that valuation. If the object in question is not a shiny crystal called ‘diamond’ but is rather a magic sword called ‘Excalibur,’ that exists only in an online game, economists would still put the value of the item at \$100,000. Similarly, if people are willing to incur large time and money costs to live in a virtual world, economists will judge that location to be lucrative real estate, regardless of the fact that it exists only in cyberspace. The mere fact that the goods and spaces are digital, and are part of something that has been given the label “game,” is irrelevant. Willingness to pay, to sacrifice time and effort, is the ultimate arbiter of significance when it comes to assessments of economic value. As avatar games consume

more human time, the assets within them will very likely grow in value; understanding how these assets are produced and traded will ultimately require a unique theory of the demand for avatar gaming.

As soon as one begins to think about an economic theory of the demand for gaming time, however, one encounters a puzzle relating to the nature of constraints. Put succinctly, in a normal market the demanders are willing to pay money to have constraints removed, but in a games market they will pay money to have constraints imposed. Think of a market for puzzles. A puzzle with nine hundred billion pieces would probably not command much willingness to pay, since the entertainment value of the game involves solving the puzzle, and that seems impossible. The agent gains emotional well-being by choosing actions that maximize the progress toward solving the puzzle, under the constraints imposed by the inherent difficulty of the puzzle. A puzzle that is too hard imposes constraints that are too severe and is no fun; relaxing the difficulty constraint should therefore raise utility and hence willingness to pay. However, a puzzle that is too easy is also no fun – who would pay money for a puzzle with only two pieces? If the puzzle went from two pieces to, say, 100 pieces, however, it would become more difficult but also more entertaining, and would therefore command a greater willingness to pay. The puzzle of puzzles is that the demand for a good can rise when a constraint becomes tighter.

The puzzle of puzzles arises primarily because economics is constructed from a model of human behavior that asserts a universal conflict between our ends and our means. The essence of behavior, to the economist, is a process of choosing actions under the constraint that we cannot have everything we want. Formally, our wants are given by

a utility function, and we seek to maximize this function subject to our constraints. If anything happens to release the constraints, say if the price of a good falls, then our utility goes up. Most economists would also assert that we are happier. Utility is good, constraints are bad. If we want to make people happier, we should remove their constraints. Hence, if we want to give people puzzles that make them happier, we should make the puzzles less challenging; by this reasoning, puzzles imposing the lightest constraints should be the most demanded in the market. This line of economic reasoning therefore leads to a deep conflict with observed behavior in game markets; players hate games that are not very challenging. It seems, then, that an economic theory of demand for puzzles, games, and other interactive entertainment goods needs to be modified in some way to allow for constraints that can raise utility and demand.

As a start toward such a theory, it is probably reasonable to assume that emotional well-being is always one goal of human behavior. People do things that make them feel happier. Secondly, it is also probably safe to assume that confronting and overcoming challenges makes people happy. Given the choice between a puzzle that is mildly challenging (put together a 100-piece puzzle) and one that is not (put together a two-piece puzzle), people will prefer the mildly challenging puzzle. At the same time, most people would prefer a 100-piece puzzle to one with 100 million pieces; the function relating challenge to fun is not monotonic. Third, if there are rewards for solving puzzles, we can assume that a puzzle with higher rewards is preferred, holding challenge levels equal.

These assumptions can be summarized in a simple economic model. Let S measure the emotional satisfaction a player receives from working on a puzzle, and let R

and C indicate the available reward and the challenge level, respectively. Then we can capture the assumptions above with a simple function like this:

$$1) \quad S = \alpha R - \beta(C - \Omega)^2$$

where Ω represents the challenge level that is ideal for the player.

Now we can introduce the utility function as a function that indicates, numerically, the intensity of an individual's desire to achieve some objective.⁹ When it comes to games and puzzles, the choice involves the amount of time to spend in one game versus another. Suppose we had games A and B, each producing satisfaction levels S_A and S_B per hour of play. Let the choosing agent have T hours to allocate between the two games. A simple utility function that illustrates the choice problem is:

$$2) \quad U(H_A, H_B) = S_A \ln(H_A) + S_B \ln(H_B)$$

where H_A and H_B are hours of play in the two games. If total time available is denoted T , hours would be allocated by maximizing the utility given in (2) subject to the constraint $T = H_A + H_B$. This setup assumes that the differing rewards and challenges of the two games produce different levels of emotional satisfaction, and that the satisfaction effects acts as weights in the motivational function. And while play time in one game does not affect the satisfaction one receives from play time in the other, there is nonetheless a diminishing marginal utility from game play: repeatedly playing the same game gets

⁹ The utility function enumerates the motivational strength of a person's goals, but it does nothing more than that. It is a numerical guide to what people will do. It is *not* a numerical guide to what makes them feel good. And while it makes sense to assume (as I do here) that the things that make people feel good will also the things that people pursue, the converse does not necessarily hold. People may well pursue things that do not make them feel good. Let S be a satisfaction ordering, such that among two states x and y , xSy implies x makes me happier than y . Let U be a motivational ordering, such that xUy implies that I am willing to pay more to obtain state x than state y . Then I think it is reasonable to assume that xSy implies xUy . However, it is not reasonable to assume that xUy implies xSy . Addiction is one counterexample. Obsession with work and money is another (Easterlin, 2001). In another paper (Castronova, 2001b), I argue that there is nothing in cultural or biological evolution that guarantees that the motivational orderings of human

boring. In this setup, the player allocates time between the games in an intuitive way: she plays games with higher rewards more often; she spends more time on games whose challenge level is not too high and not too low; and she will play a game that is less inherently satisfying, at least for a time, simply for the variety of it.¹⁰

We can introduce the price of gaming as follows. Let p_A and p_B be the prices of games A and B respectively, and let G represent consumption of all goods other than game play.¹¹ Let $p_G = 1$. The utility function will have G as a third argument, but if the agent has Y dollars of income to spend on games and other goods, then we have

$$3) \quad G = Y - p_A H_A - p_B H_B$$

The objective function regulating hours of game time could then be expressed as

$$4) \quad U(H_A, H_B, Y - p_A H_A - p_B H_B) = S_A \ln(H_A) + S_B \ln(H_B) + \gamma \ln(Y - p_A H_A - p_B H_B)$$

As above, the agent would solve this problem to find the optimal levels of H_A and H_B , and would allocate time accordingly. Games that are more fun would be played more often; no one would devote all of his time to one game; games that are more expensive to play would be played less often.

Whether this is a particularly elegant approach to the market for games is open to debate, of course, but it certainly is useful for exploring some of the unusual features of

beings must also be their satisfaction orderings. On the contrary, the process of economic development introduces a systematic bias that points our motivational orderings *away* from our satisfaction orderings.

¹⁰ One implication of this framework that is worth exploring: a competitive market in games will generally not produce games with the ideal challenge level for a given player. Game content is costly to produce, and maintaining both rewards (R) and the challenge level (C) can only be done at some marginal cost. As long as the marginal cost of challenge is positive, competitive game companies will introduce challenge to the extent that its marginal cost of production equals its marginal revenue to the company. The marginal revenue must therefore be positive, which would seem to imply that the marginal contribution of challenge to emotional satisfaction must be positive. This will only happen at values of C that are below Ω . In competitive markets for puzzles, all puzzles will be at least a little bit too easy.

¹¹ Pricing in MMORPGs is primarily a flat-rate monthly fee for game access. Therefore, in the model, I ignore the one-time cost of buying the game software and concentrate instead on the idea that gaming must be purchased in units of time.

that market, specifically its interaction with real world labor markets. According to an earlier paper on EverQuest (Castronova, 2001a), many people spend more time in games like EverQuest than they do at work. Moreover, those who devoted more time to the game seem to have somewhat lower wage rates, but not dramatically lower. At the same time, that paper documented the fact that people can make real money by selling the digital items that they produce while playing. There seems to be a distinct emigration of work time from Earth to Norrath. This is an important aspect of real world gaming, and it has distinct, and odd, implications in the context of this choice model.

To explore this further, assume the choice is between work time, denoted by L , and time in a single game, denoted by H . Let game time have a price p , and produce satisfaction S , and let work be compensated at the wage rate w . Then

$$5) \quad Y = wL$$

Let non-gaming leisure time be Z , a third variable in the utility function, given by

$$6) \quad Z = T - H - L$$

Continuing with a log-linear utility function, we have:

$$7) \quad U(H, wL - pH, T - H - L) = S \ln(H) + \gamma \ln(wL - pH) + \delta \ln(T - H - L)$$

The agent allocates time among the game and work so as to maximize Equation (7).

While the problem does not admit a simple solution, inspection of it reveals a number of interesting features of the demand for games. First, the constraint aspect of the games, given by the challenge level C , enters the model as a weight on the utility function. This separates it mathematically from the constraints of time and money, which, in most utility maximization problems, appear as constraint equations and only enter utility indirectly, as in (7). Conceptually, this allows us to think of two kinds of

constraints in the world: the traditional constraints of economics (time and money), whose relief always results in higher utility, and the new category of gaming constraints, whose relief may actually reduce emotional satisfaction, and hence result in lower utility. The Puzzle of Puzzles is resolved here by recognizing that constraints can have a positive effect on emotional satisfaction, and therefore states with tougher constraints may actually be more desirable. The utility function is then designed to give higher weight to more desirable states, since they are preferred over less desirable states. As a result, utility can be higher when a constraint is tougher; hence there can be a willingness to pay for tougher constraints.

A second aspect worth mentioning is the possibility that wages have both income and substitution effects with respect to game time. People with higher wages tend to be richer, hence they will demand more of all normal goods, including game time. However, highly-paid people also face a higher opportunity cost of gaming, hence they will demand less. This suggests that game demand may be U-shaped with respect to wages. Very well-paid people can play more because they can afford all kinds of leisure activities. Poorly-paid people can play more because they are not sacrificing very much income to do so. Conversely, those with moderate wages may be very sensitive to the impact of gaming time on their earnings and careers.

Third, note that money enters this problem in an unconventional way. As usual, we have goods prices in the budget constraint, and wages as the price of leisure. What is new is the possibility that money can enter the problem as a parameter of the utility function. This happens if the rewards of playing (R), which affect gaming satisfaction (S), happen to be partly denominated as cash. As mentioned above, in a game like

EverQuest, players can make substantial amounts of money by farming the virtual world and selling the produce in internet truck markets like Ebay. Some of these players explicitly consider these funds as income; some think of their farming as a job, as work, not play. How should economists approach this? Perhaps it is a sub-problem, where the agent must choose to allocate time between Earth work, virtual work, and virtual play. Or, perhaps these game earnings should be treated as a price discrimination scheme, effectively lowering the net price of the game for the more serious players. In that case, farm receipts should be taken out of R , and instead be subtracted from p in the budget constraint. A third approach would be more radical: instead of thinking of game time as partly work, perhaps we should think of work time as just another game. Then the issue can be handled elegantly in Equation (2). Game A happens to be the always-exciting “Work Game of Earth,” where you go to the office and face the challenges, denoted C , presented by your boss, your co-workers, and your competitors, and where overcoming those challenges garners you rewards, denoted R , in the form of wages, perks, fringe benefits, and assorted entertainments involving the office copy machine.¹² People who get more satisfaction from Game A will put more time into it. Nonetheless, Game A can get boring, so even the most rabid fan of Game A will be observed putting some time into Game B. Regardless of how it is approached, it is clear that there is a substitution between earth work and game time that depends, to some extent, on the financial rewards available in each.

Thus, simple as it is, the framework developed so far throws light on the two most critical aspects of gaming as an emerging economic phenomenon: Game time is a

¹² One way of judging whether the Work Game of Earth is entertaining is by counting how many people play it. By that standard, it seems to be a very entertaining game. Indeed, it seems to be especially fun for

substitute for other consumption goods, and it is also a substitute for work time. The degree to which this substitution occurs depends on wages and prices both on earth and inside the games. It also depends on the emotional satisfactions and general costs of game time. In the most radical approach to game/work substitution, the emotional satisfaction of Earth work are directly compared to the emotional satisfaction of game time. These Earth/game substitutions involve real economic transfers. It has already been shown (Castronova, 2001a) that labor devoted to games produces durable economic assets with observable market values. The wealth stock and annual production of a game world is already significant on a per capita basis; they will become significant macroeconomic aggregates if the stream of earth to game substitution becomes strong. And we can gauge the potential strength of substitutions into gaming by asking how satisfying and costly game time may possibly become in the future.

IV. The Market for Virtual Worlds: Technology and Market Structure

We cannot see the future, of course, but there are a number of technological innovations that are relevant to gaming, that are also fairly easy to see coming. Currently, access to gaming involves some sort of access to computing technology, and access to gaming that can earn money involves access to a shared, persistent, physical computing environment, specifically a virtual world (Castronova, 2001a). The technology supporting virtual worlds is advancing so quickly that it would be foolish to describe the next generation in any detail. Suffice it to say that there are large, lucrative industries working energetically on different dimensions of the environment that virtual worlds thrive in.

people in richer cultures. Cause and effect are hard to sort out in this case, however.

These industries produce three items of interest, namely, Connection, Interface, and Content. Developments in connections involve the internet and, increasingly, wireless communications. Development of interfaces involves voice command, heads-up displays, and body motion detection (computer-controlling gloves, gaze readers). Developments in content involve the supply side of the market for games, where annual revenues have grown beyond Hollywood box office revenues.¹³ All three industries are expanding at a rapid rate. Whatever emotional experiences people seek, it may become possible, in the near future, to effortlessly connect to a virtual world that provides that experience at fairly low cost. Kurzweil (1999) argues that the explosion of computing power alone may be sufficient to change the daily course of life.

Since these developments all involve networks, they may seem to suggest a monopolistic market structure. If economic life online involves getting your email and hanging around with friends, there will be positive externalities with respect to the sheer size of the virtual world one visits. If I spend my time on Rubi-Ka, while you spend your time in Albion, we cannot talk to one another, and we cannot do things together. Thus, our time in virtual worlds is more valuable if everyone we know is in the same world. Moreover, if two worlds compete and one has more players than another, wouldn't everyone have an incentive to join the larger world, so as to enjoy the larger network of society, communication, and entertainment that it affords? Might such network externalities lead to a domination of this market by one player?

There are reasons to expect, however, that this market is not likely to be monopolized. First, there seems to be a great diversity of tastes for the different features

¹³ Perhaps the Turing Test will first be passed in a game. Content developers have been focusing especially hard on the artificial intelligence of software agents, with some success (Johnson, 2002). University of

of a world. Mr. Bird may want to be on Pluto, while Mr. Castronova prefers medieval Britain. One of the major attractions of life mediated by avatars is the anonymity it affords, and anonymity requires a person to have exit options, other worlds to escape to if one's reputation in this one gets unpleasant. Perhaps a savvy game developer could make a world so large and varied as to provide the essential minimum level of entertainment and anonymity to a sufficiently large number of people, so that membership in that one world becomes optimal for all. This seems unlikely, however, given that there is a marginal cost to creating and maintaining game content.¹⁴ Moreover, there are no economies of scale on the supply side to match the increasing returns on the demand side (Liebowitz and Margolis, 1994). Production of game content and its maintenance are both labor intensive activities. One could perhaps increase production of content by allowing other producers (say, by opening your game code to the public), but continued control of the world being created would be problematic.¹⁵ On the whole, it seems very unlikely that one developer could produce a world big enough to monopolize the market.

A second reason involves congestion. Virtual worlds are virtual because they are online, but they are worlds because there is some physicality to them. Avatars take up space. If a world has a certain amount of entertaining content in it, that content will almost always be subject to some kind of congestion effect. The cool monsters are in the Dungeon of Befallen, but if tens of thousands of us go there to hunt them, none of us will have a good time. Often the only way to reduce congestion is to add content, which,

Michigan computer scientist John Laird has argued that games represent the forefront of AI research.

¹⁴ At the 2002 Electronic Entertainment Expo, developers at a workshop on virtual worlds repeatedly insisted that the ongoing customer support costs required to keep the world in existence equaled or surpassed the entire development cost.

¹⁵ This is the strategy of Project Entropia (Mindark) and Neverwinter Nights (Bioware).

again, is labor intensive. There will also be congestion effects related to connection speeds and bandwidth.

A third reason that the market will probably not be dominated by a few companies can be found in the many competitive strategies that are available even now, but have not yet been exploited by new entrants. For example, the current set of developers have managed to impose huge switching costs on players by structuring game play around the time-intensive development of avatar capital. A player starts the game with a weak avatar, but game play gives the avatar ever-increasing powers. As power increases, the avatar is able to take more advantage of the game world, to travel farther, do more things, see more people. A person with a high-level avatar then faces a high hurdle in switching games, because in the new game he will start out poor, defenseless, and alone again. This situation definitely locks in the game's player base, but it is also open to defeat by any number of schemes to reduce the switching costs. Surprisingly, no competitor to a current game has offered new players the opportunity to start their avatars at a higher level of wealth and ability if they can provide evidence of a high level avatar in another game. Nor have they offered the opportunity to buy their way to a powerful avatar from the start. In fact, by not offering this kind of monetary opt-in, the companies implicitly encourage the buying and selling of avatars outside the game. In sum, what appear to be strong lock-ins and switching costs in the game market today may not be as strong as they seem; when savvy competitors appear, the player bases will generally be at risk.

A final argument against a monopolization tendency comes from the nature of the content itself. Games are art, for the most part, and markets for artistic output exhibit a great deal of churn due to herding effects and the star phenomenon (MacDonald, 1988).

If a company designs a better game, it will attract players. And while it is true that development costs can be significant, it will always be possible to produce a fun virtual world for a tiny amount of money and then scale it up as it becomes more popular.¹⁶

Whatever network externalities, supply-side returns to scale, and barriers to entry may exist in the market for virtual worlds, they seem insufficient to produce domination by a single company. The distribution of populations in virtual worlds is perhaps less like a natural monopoly market than a club goods market. Populations will sort according to the services, ambience, and fees of the various worlds. Virtual worlds will compete, as clubs do, but their size will be limited by congestion effects and by the marginal cost of increasing the scale of the world.

Having considered the historical technology dependence of avatar games, the fairly unusual demand theory that they require, and what seems to be their most natural market structure, we are in a position to shed some light on the first of the two questions that have been the object of the study: What will the online multiplayer gaming market look like in the medium-term future? The preceding sections suggest that it will consist of a number of large, densely populated worlds, with varying degrees of portability between them. These worlds will generate large revenue streams and will occupy many hours of human time, some of it considered play, some of it considered work. If current behavioral patterns continue, many of the hours that people devote to games will result in the accumulation of stocks of digital capital goods. These objects will have real economic value, and the total capital hosted by online worlds will be proportionate to the amount of

¹⁶ The developers at E3, mentioned above, also asserted that no virtual world could be developed for less than \$15 million. Nonetheless, Mythic Entertainment developed the successful world of Dark Age of Camelot for only \$2.5 million. It is no coincidence that the workshop speakers were representatives from

total hours of labor input the worlds receive. Given the expected growth in connectivity, interface technologies, and content, there is reason to believe that the digital capital stock may eventually become quite large.

These considerations then lead to the next set of questions: If virtual worlds do become more important, how will this affect the real Earth economy?

V. A Macroeconomic Impact of Virtual Economies?

If virtual worlds do, in fact, grow as a human phenomenon, there may be some implications for Earth economies. The implications are easiest to see if we ignore for a moment the fact that these are online games, and just imagine what would happen if humans became obsessed with a general time-consuming hobby, or perhaps a hallucinogenic drug, that, from the perspective of the Earth economy, had similar effects on behavior. The effects of most interest are:

- 1) People pursuing this hobby/taking this drug generally spend 60 to 80 hours a week in an activity that has essentially no impact, positive or negative, on the economy. They simply check out of Earth.
- 2) Even when they are not pursuing the hobby/taking the drug, these folks have very little interest in either advancement, or entertainment, in Earth society. They are able to get some Earth cash by selling trinkets that they make while hobbying. But their hearts are not in their Earth earnings. Indeed, they come to Earth only to work a little bit, and get just enough Earth money to support

the very large players in this market (Sony, Universal/Vivendi, Microsoft), who of course had every incentive to quash expectations of success among the many tiny competitors in the audience.

their Earth needs, such as food, water, some simple clothes, and a roof over their head.

These behaviors are hypothetical, of course, but they are credible in so far as they represent entirely normal behavior on the part of a substantial fraction of the players of games like EverQuest (Castronova, 2001a). The question is, what would happen if this kind of behavior became much more widespread?

First, it is important to recognize that this change does not necessarily mean that people are worse off. The mere fact that labor hours that were once producing automobiles are now producing avatars does not mean anything about the level of wealth in society. The basket of produced goods is simply changing. Players will be creating new assets, with real economic value, in online worlds. A proper accounting could very easily show that the actual production of wealth per capita is rising.

The difficulty is that current national income and product accounts do not place any value on online assets. Nor do they seem likely to do so anytime soon. Concepts like the GDP are nation-based, but to what Earth nation do assets in virtual worlds belong? The answer is “none,” at the moment anyway. As a result, a migration of value creation from Earth economies to virtual economies would appear as a decline in standard measures of economic activity, such as the GDP. Earth economies would seem to be in recessions or depressions.

A second impact involves the demographic structure of the transition. If devotion of time to virtual economies has a U-shaped relationship to the wage, we might predict that migration to virtual worlds would proceed much like the usual Earth migrations. The vast majority of émigrés from Earth would be those whose wages on Earth are low

relative to their wages in cyberspace. And then there would be a substantial number of very well-paid people for whom interworld travel is relatively costless. Together, both groups might represent a significant brain drain from the earth economy, the former group because tech savvy is probably going to receive a higher return in cyberspace than on Earth, the latter because the well-paid may find virtual worlds generally more entertaining than Earth. On the other hand, the opening of a cyberian frontier, like other frontiers, will have a very refreshing leveling effect: those whose Earth shapes expose them to brutality, stigma, and insufferable limitations will find freedom and relief when they live through less stigmatized virtual shapes.

A final impact worth noting involves the fiscal health of earth governments. If economic activity migrates into virtual economies, where there are no earth jurisdictions, there will be a net loss of taxable assets and incomes in Earth economies. At the same time, there may be substantial reductions in demand for Earth government services (i.e. roads). There may be long periods of time in which the tax base is eroding more rapidly than service demands, and there may be severe inequalities across jurisdictions in these rates as well.

Taking these fiscal policy effects together with the possible labor supply and GDP shocks, it would seem that a large migration to the cyberian frontier could conceivably impose serious stress on Earth political systems. Whether or not these shocks and stresses actually appear depends on the degree to which the connection, interface, and content industries succeed in their efforts to produce immersive gaming experiences on a massive scale.

Of course it is not possible to see specific aspects of the future with much accuracy. What seems most likely, on a broader level, is that a large migration of economic activity into cyberspace would have to have some impact on the way that one conceives of the macroeconomy. New statistics and economic management policies may have to be developed. However, if the emergence of virtual worlds does eventually require some governmental reaction, it is still not clear which Earth governments should be involved. Virtual worlds seem to exist as separate political entities at the moment, and this raises new constitutional issues.

VI. Constitution and Governance

Indeed, the most salient current policy issues both within and outside of games involve issues of governance. In the US, there have been judicial rulings indicating that Earth courts have no jurisdiction over events that occur online (Kaplan, 2001). An argument (unsuccessful in court) has been made that video games are speech, and are therefore entitled to constitutional protections that would make game companies the de facto legitimate governments of their game worlds (Au, 2002). However, players in these games are citizens of Earth countries and their incomes from game activities are certainly subject to tax. A full-blown legal analysis would be inappropriate here, but it is evident that there is confusion about who is ultimately responsible for events that occur in virtual worlds. As the value of virtual world assets and trade rises, economic agents will have ever greater incentives to seek the usual protections, damages, and claims from some higher authority. Moreover, the real emotional investment of people in their online lives will almost certainly lead them to seek out a forum where their grievances may be aired

and then acted upon with force (Becker, 2002). Only time will tell who the governing authority will ultimately be.

Earth courts may eventually be the final authority, and Earth governments may be another. But at the moment, the game owners are effectively filling this role, with interesting implications. Their power derives from the fact that every player who logs on to a game accepts an End User Licensing Agreement (EULA) that strongly limits their rights to affect events in the game world. Under Sony's EverQuest EULA, every click and motion in the game is defined as 'uploaded content,' to which the player waives any and all rights of control. A player could therefore develop in-game assets worth hundreds of thousands of dollars, have those assets wiped out by a coding error, and have no recourse for damages. To handle such issues, the game companies put significant resources into their customer service operations. Mythic Entertainment (developer of Dark Age of Camelot) formally appointed a "Player Representative" to act as a customer service spokesperson. She expressed the state of affairs clearly in a discussion board post:

"Any one issue might have several viewpoints, all of which are probably represented within the company itself. I can understand the frustration that people feel when they don't hear anything about their pet issue (because I feel it myself), but the fact is it doesn't get discussed publicly until a decision has been made. Games are not democracies. SOMEONE has to drive, and as I've said on several boards, I as the player representative do not drive the Good Ship Mythic, I am merely the most annoying backseat driver ever. The only 'votes' are called dollars. If you aren't having fun, you shouldn't be playing." (Sanya Thomas, at player2player.net/forums/ on June 6, 2002.)

Evidently, game owners are dictators whose benevolence depends only on the constraint that they must remain profitable.

This power structure has predictable effects. In every game currently on the market, the owners consider it their right to introduce changes to game mechanics at any time, without prior consultation with the players. As a result, avatars can have their real market value destroyed overnight, without warning. The only option for players is to complain loudly at various fansite discussion boards, and the players make use of this privilege zealously. A typical board (i.e. eqvault.ign.com) is flooded every day with arguments, suggestions, and pleas, of a breathtakingly varying quality, about every aspect of the game. Any change to the game is immediately met with howls of protest from those damaged most; those who gain, typically, say nothing because they are in the game, enjoying their new benefits. Game owners occasionally seem to pay attention to these forums, which must represent only the tip of the iceberg of player input in the form of emails and other communications. The net result is that the political structure of every virtual world consists of a group of all-powerful executives surrounded by mobs of angry, harassing supplicants.

As an example of ongoing governance problems, consider foreign trade policy, currently one of the most pressing issues facing this polity. By 'foreign trade' is meant the common practice of selling in-game items for real money in out-of-game markets. This trade is simple to conduct and hard to detect. If Castronova has an avatar that owns a Flowing Black Silk Sash in EverQuest, he may go to his neighbor Bird and sell the sash to him for \$100. Bird simply gives \$100 to Castronova, and then both return to their computer rooms, launch their avatars in EverQuest, and meet at some pre-assigned place in the virtual world. Once there, Castronova's avatar gives Bird's avatar the sash, completing the transaction.

This practice puts game owners in a quandary. On the one hand, all transactions like this improve the well-being of both parties, and therefore make their enjoyment of the game greater. They are happier customers. On the other hand, widespread foreign trade can ruin the ambience of the game world. Most games seek to give the player a rags-to-riches experience, but the satisfaction of that experience can be significantly lessened if one observes that other players, who ought to be poor like oneself, are instead very well arrayed in expensive equipment that they bought for hundreds of dollars outside the game. Foreign trade therefore erodes the equality of opportunity of game play, and damages the entire gaming environment; the situation is a commons tragedy, where the self-interested trading behavior of individuals destroys the game's atmosphere, to the detriment of all. Whether or not to allow foreign trade therefore involves deep questions about the purpose of the game, the desired atmosphere, and the interests (economic and emotional) of all the players.

The game companies have taken varying stances, from formally outlawing the practice, with and without serious enforcement efforts, to complete laissez faire, and policy pronouncements in this arena have had dramatic effects on the value of assets and the quality of the gaming atmosphere. One company's efforts to control foreign trade did produce a wonderful gaming atmosphere, but resulted in a formal court action by market-oriented players (Becker 2002). Without taking a position on this and other cases, one thing is clear: foreign trade policy has certainly been imposed *on* the people rather than *with* the people.

To anyone versed in political history, it should be no surprise that the game companies have made themselves vulnerable by approaching these matters as customer

service issues rather than governance. In their own minds, the players are not customers, but citizens, with corresponding rights. Indeed, “A Declaration of the Rights of Avatars” has already been proclaimed (see Raph Koster’s work at www.legendmud.org/raph/gaming/index.html). Little wonder, then, that player-company relations tend to be very tense, even in the best games.¹⁷ There seems to be some possibility that game company autarchs may follow Frederick the Great into the dustbin of history. The customer service state, like all benevolent despotisms, suffers from illegitimacy.

On the other hand, unlike Frederick the Great, a game company must make decisions that meet the profit test. And while the players may be powerless within the game, they are not serfs. They have both voice and exit as options for resistance. Thus, survival in the competitive world of gaming requires that a company remain popular with its gamers. Ironically, increasing freedom of movement of labor and capital on earth has already forced earth governments to compete with one another to induce people to remain under their jurisdiction. A similar process seems to be occurring with game governments, where cross-jurisdictional mobility is already quite liquid. Perhaps the future will see earth governments and game governments in competition with one another. A Darwinian struggle for jurisdictional reach might provide an outstanding system of checks and balances, especially considering the ease of movement in cyberspace.

The net results of this jurisdictional competition are very hard to predict. It seems most likely that populations will sort according to tastes, with those who desire some

¹⁷ Mythic Entertainment has made a very serious effort to remain engaged with the player base of Dark Age of Camelot, but they have been unable to avoid the basic pattern described above. The loss of faith in the company does not stem from bad decisions, but from the impossibility of endowing any kind of dictatorial decisions - even good ones - with legitimacy.

voice in their affairs seeking more democratic forms of game governance. Those who want a non-market, equality-of-opportunity game world will be able to seek that out; those who wish to buy and sell their way to the top will find an arena that suits those tastes. Overall, however, it seems likely that constitutional issues will be important for some time.

VII. Conclusion: On the Uniqueness of Virtual Economies

This paper has attempted to describe some of the unique features of economies in virtual worlds. Living in these worlds involves a leap into a fantasy existence, something that humans have apparently been trying to do since the dawn of civilization. The demand for game time can be expressed in a simple economic model, and it seems to have increased as the immersive satisfaction available from gaming technology has increased. If this pattern continues, the advances of the information age could make gaming a significant aspect of the lives of millions of people. That scenario may have macroeconomic implications, as well as some effects on government policy.

A common theme throughout the paper is that the analysis of virtual economies will require slightly different tools and approaches than we are used to. The differences are dictated by the specific features of life in cyberspace. In virtual worlds, the entire physical universe is open to direct and costless manipulation by the owners of the game. The human beings behind the avatars are real, and physical, and subject to the laws of Earth, but the avatars themselves do not inherently face any physical constraints at all. The discovery and description of avatar-mediated economic life represent the most important current research avenues in the economics of games.

Indeed, further thinking about some of the topics in the preceding sections reveals a number of areas in which the behaviors and outcomes that we generally take to be standard in Earth economics do not seem to hold in avatar economies. Some examples:

- Economics, on Earth, argues that no wise government will try to control prices. In an avatar economy, however, the government can effortlessly peg many prices at any value. Since the goods are digital, they can be costlessly created and

destroyed. Hence price ceilings create no excess demand, and price floors no excess supply. It may make sense to control some prices.

- Economics, on Earth, assumes that work causes disutility. In an avatar economy, however, it is lack of work that causes disutility. Regardless of earnings and loot rates, people who play games must have something to do or they will be bored. If a game structure limits their ability to be meaningfully engaged in some mission, quest, or activity, they will be unhappy. Work is good.
- Economics, on Earth, believes that economic growth is always good. In an avatar economy, however, increases per-capita wealth – which make it easier to accomplish various quests and missions – will lower the challenge level of the game, potentially making it a less interesting puzzle. Growth can be bad.
- Economics, on Earth, takes the population of humans as fixed, and also assumes that their tastes and initial abilities are fixed. In an avatar economy, however, people are free to choose a significant subset of their abilities. They also can choose when to be alive and when not to be, as well as how many different people to be. The choosing economic agent can be a fairly complex entity.

These examples present a number of puzzles for economic research. It should be possible to generate fairly simple theories and arguments explaining why things do seem somewhat different in virtual economies than they do in the Earth economy. As those arguments are made, we will learn more about the things that are the same in all economies, both virtual and Earthly: the true nature of human motivation and well-being, and their true relationship to objects in the immediate physical world.

References

- Au, Wagner James (2002), "Playing Games With Free Speech," *Technology and Business*, salon.com, May 6.
- Becker, David (2002), "Game Exchange Suit Goes to Court," *CNET News.com*, February 7.
- Borges, Jorge Luis (1962), "The Lottery in Babylon," translated by John M. Fein, in Borges, *Labyrinths: Selected Stories and Other Writings*, New York: New Directions Publishing.
- Castronova, Edward (2001a), "Virtual Worlds: A First-Hand Account of Market and Society on the Cyberian Frontier," CESifo Working Paper No. 618, December.
- _____. (2001b), "Achievement Bias in the Evolution of Preferences," Gruter Institute Working Papers on Law, Economics, and Evolutionary Biology, Volume 2.
- Easterlin, Richard A. (2001), "Income and Happiness: Towards a Unified Theory," *Economic Journal*, 111, 465-484.
- Eco, Umberto (1989), *Foucault's Pendulum*, translated by William Weaver, New York: Harcourt Brace Jovanovich.
- Engwall, Lars (1994), "Bridge, Poker, and Banking," in Donald E. Fair and Robert Raymond, eds., *The Competitiveness of Financial Institutions and Centres in Europe*, Dordrecht: Kluwer Academic Publishers, 227-39.
- Johnson, Steven (2002), "Wild Things," *Wired*, March.
- Kaplan, Carl S. (2001), "Florida Community Can't Shut Down 'Voyeur Dorm,'" *New York Times*, October 5.
- Kurzweil, Ray (1999), *The Age of Spiritual Machines*, New York: Penguin Books.
- Liebowitz, S.J. and Stephen Margolis (1994), "Network Externality: An Uncommon Tragedy," *Journal of Economic Perspectives* 8(2), 133-150.
- MacDonald, Glenn M. (1988), "The Economics of Rising Stars," *American Economic Review*, 78(1), 155-66.
- Metrick, Andrew (1995), "A Natural Experiment in 'Jeopardy!'" *American Economic Review*, 85(1), 240-53.
- Mazalov, Vladimir V., Svetlana V. Panova, and Mojca Piskuric (1999), "Two-Person Bilateral Many-Rounds Poker," *Mathematical Methods of Operations Research* 49(2), 267-82.
- Nash, John F. and Shapley, L. S. (1997), "A Simple Three-Person Poker Game," in Mary Ann Dimand and Robert W. Dimand, eds., *The Foundations of Game Theory, Volume II*, Cheltenham: Elgar Reference Collection, 13-24.
- Nichols, Mark W. (1998), "The Impact of Deregulation on Casino Win in Atlantic City," *Review of Industrial Organization*, 13(6), 713-26.
- Page, Scott E. (1998), "Let's Make A Deal," *Economics Letters* 61(2), 175-80.
- Shubik, Martin (1999), "The 'Bridge Game' Economy: An Example of Indivisibilities," in Martin Shubik, ed., *The Selected Essays of Martin Shubik*, Cheltenham: Elgar, 184-187.
- Simon, Herbert A., and Jonathan Schaeffer (1992), "The Game of Chess," in Robert J. Aumann and Sergiu Hart, eds., *Handbook of Game Theory With Economic Applications*, Amsterdam: Elsevier, 1-17.

CESifo Working Paper Series

- 686 Hyun Park and Apostolis Philippopoulos, Can Productive Government Spending be the Engine of Long-Run Growth When Labor Supply is Endogenous?, March 2002
- 687 Jonathan P. Thomas and Tim Worrall, Gift-Giving, Quasi-Credit and Reciprocity, March 2002
- 688 Barbara Buchner, Carlo Carraro, Igor Cersosimo, and Carmen Marchiori, Back to Kyoto? US Participation and the Linkage between R&D and Climate Cooperation, March 2002
- 689 Amihai Glazer and Vesa Kanninen, The Effects of Employment Protection on the Choice of Risky Projects, March 2002
- 690 Michael Funke and Annkatrin Niebuhr, Threshold Effects and Regional Economic Growth – Evidence from West Germany, March 2002
- 691 George Economides, Apostolis Philippopoulos, and Simon Price, Elections, Fiscal Policy and Growth: Revisiting the Mechanism, March 2002
- 692 Amihai Glazer, Vesa Kanninen, and Mikko Mustonen, Innovation of Network Goods: A Non-Innovating Firm Will Gain, March 2002
- 693 Helmuth Cremer, Jean-Marie Lozachmeur, and Pierre Pestieau, Social Security, Retirement Age and Optimal Income Taxation, April 2002
- 694 Rafael Lalive and Josef Zweimüller, Benefit Entitlement and the Labor Market: Evidence from a Large-Scale Policy Change, April 2002
- 695 Hans Gersbach, Financial Intermediation and the Creation of Macroeconomic Risks, April 2002
- 696 James M. Malcomson, James W. Maw, and Barry McCormick, General Training by Firms, Apprentice Contracts, and Public Policy, April 2002
- 697 Simon Gächter and Arno Riedl, Moral Property Rights in Bargaining, April 2002
- 698 Kai A. Konrad, Investment in the Absence of Property Rights: The Role of Incumbency Advantages, April 2002
- 699 Campbell Leith and Jim Malley, Estimated General Equilibrium Models for the Evaluation of Monetary Policy in the US and Europe, April 2002
- 700 Yin-Wong Cheung and Jude Yuen, Effects of U.S. Inflation on Hong Kong and Singapore, April 2002

- 701 Henry Tulkens, On Cooperation in Musgravian Models of Externalities within a Federation, April 2002
- 702 Ralph Chami and Gregory D. Hess, For Better or For Worse? State-Level Marital Formation and Risk Sharing, April 2002
- 703 Fredrik Andersson and Kai A. Konrad, Human Capital Investment and Globalization in Extortionary States, April 2002
- 704 Antonis Adam and Thomas Moutos, The Political Economy of EU Enlargement: Or, Why Japan is not a Candidate Country?, April 2002
- 705 Daniel Gros and Carsten Hefeker, Common Monetary Policy with Asymmetric Shocks, April 2002
- 706 Dirk Kiesewetter and Rainer Niemann, Neutral and Equitable Taxation of Pensions as Capital Income, April 2002
- 707 Robert S. Chirinko, Corporate Taxation, Capital Formation, and the Substitution Elasticity between Labor and Capital, April 2002
- 708 Frode Meland and Gaute Torsvik, Structural Adjustment and Endogenous Worker Recall Probabilities, April 2002
- 709 Rainer Niemann and Caren Sureth, Taxation under Uncertainty – Problems of Dynamic Programming and Contingent Claims Analysis in Real Option Theory, April 2002
- 710 Thomas Moutos and William Scarth, Technical Change and Unemployment: Policy Responses and Distributional Considerations, April 2002
- 711 Günther Rehme, (Re-)Distribution of Personal Incomes, Education and Economic Performance Across Countries, April 2002
- 712 Thorvaldur Gylfason and Gylfi Zoega, Inequality and Economic Growth: Do Natural Resources Matter?, April 2002
- 713 Wolfgang Leininger, Contests over Public Goods: Evolutionary Stability and the Free-Rider Problem, April 2002
- 714 Ernst Fehr and Armin Falk, Psychological Foundations of Incentives, April 2002
- 715 Giorgio Brunello, Maria Laura Parisi, and Daniela Sonedda, Labor Taxes and Wages: Evidence from Italy, May 2002
- 716 Marta Aloi and Huw Dixon, Entry Dynamics, Capacity Utilisation and Productivity in a Dynamic Open Economy, May 2002
- 717 Paolo M. Panteghini, Asymmetric Taxation under Incremental and Sequential Investment, May 2002

- 718 Ben J. Heijdra, Christian Keuschnigg, and Wilhelm Kohler, Eastern Enlargement of the EU: Jobs, Investment and Welfare in Present Member Countries, May 2002
- 719 Tapio Palokangas, The Political Economy of Collective Bargaining, May 2002
- 720 Gilles Saint-Paul, Some Evolutionary Foundations for Price Level Rigidity, May 2002
- 721 Giorgio Brunello and Daniela Sonedda, Labor Tax Progressivity, Wage Determination, and the Relative Wage Effect, May 2002
- 722 Eric van Damme, The Dutch UMTS-Auction, May 2002
- 723 Paolo M. Panteghini, Endogenous Timing and the Taxation of Discrete Investment Choices, May 2002
- 724 Achim Wambach, Collusion in Beauty Contests, May 2002
- 725 Dominique Demougin and Claude Fluet, Preponderance of Evidence, May 2002
- 726 Gilles Saint-Paul, Growth Effects of Non Proprietary Innovation, May 2002
- 727 Subir Bose, Gerhard O. Orosel, and Lise Vesterlund, Optimal Pricing and Endogenous Herding, May 2002
- 728 Erik Leertouwer and Jakob de Haan, How to Use Indicators for ‘Corporatism’ in Empirical Applications, May 2002
- 729 Matthias Wrede, Small States, Large Unitary States and Federations, May 2002
- 730 Christian Schultz, Transparency and Tacit Collusion in a Differentiated Market, May 2002
- 731 Volker Grossmann, Income Inequality, Voting Over the Size of Public Consumption, and Growth, May 2002
- 732 Yu-Fu Chen and Michael Funke, Working Time and Employment under Uncertainty, May 2002
- 733 Kjell Erik Lommerud, Odd Rune Straume, and Lars Sørgard, Downstream Merger with Oligopolistic Input Suppliers, May 2002
- 734 Saku Aura, Does the Balance of Power Within a Family Matter? The Case of the Retirement Equity Act, May 2002
- 735 Sandro Brusco and Fausto Panunzi, Reallocation of Corporate Resources and Managerial Incentives in Internal Capital Markets, May 2002
- 736 Stefan Napel and Mika Widgrén, Strategic Power Revisited, May 2002
- 737 Martin W. Cripps, Godfrey Keller, and Sven Rady, Strategic Experimentation: The Case of Poisson Bandits, May 2002

- 738 Pierre André Chiappori and Bernard Salanié, Testing Contract Theory: A Survey of Some Recent Work, June 2002
- 739 Robert J. Gary-Bobo and Sophie Larribeau, A Structural Econometric Model of Price Discrimination in the Mortgage Lending Industry, June 2002
- 740 Laurent Linnemer, When Backward Integration by a Dominant Firm Improves Welfare, June 2002
- 741 Gebhard Kirchgässner and Friedrich Schneider, On the Political Economy of Environmental Policy, June 2002
- 742 Christian Keuschnigg and Soren Bo Nielsen, Start-ups, Venture Capitalists, and the Capital Gains Tax, June 2002
- 743 Robert Fenge, Silke Uebelmesser, and Martin Werding, Second-best Properties of Implicit Social Security Taxes: Theory and Evidence, June 2002
- 744 Wendell Fleming and Jerome Stein, Stochastic Optimal Control, International Finance and Debt, June 2002
- 745 Gene M. Grossman, The Distribution of Talent and the Pattern and Consequences of International Trade, June 2002
- 746 Oleksiy Ivaschenko, Growth and Inequality: Evidence from Transitional Economies, June 2002
- 747 Burkhard Heer, Should Unemployment Benefits be Related to Previous Earnings?, July 2002
- 748 Bas van Aarle, Giovanni Di Bartolomeo, Jacob Engwerda, and Joseph Plasmans, Staying Together or Breaking Apart: Policy-makers' Endogenous Coalitions Formation in the European Economic and Monetary Union, July 2002
- 749 Hans Gersbach, Democratic Mechanisms: Double Majority Rules and Flexible Agenda Costs, July 2002
- 750 Bruno S. Frey and Stephan Meier, Pro-Social Behavior, Reciprocity or Both?, July 2002
- 751 Jonas Agell and Helge Bennismarker, Wage Policy and Endogenous Wage Rigidity: A Representative View From the Inside, July 2002
- 752 Edward Castronova, On Virtual Economics, July 2002