Game Studies

the international journal of computer game research

volume 3, issue 2
december 2003

Edward Castronova is Associate Professor of Economics at California State University, Fullerton, USA.

Author's homepage

On Virtual Economies

by Edward Castronova

Abstract:

Currently, several million people have accounts in massively multiplayer online games. The population of virtual worlds has grown rapidly since 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 simple choice model of the demand for game time. 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.

I. Introduction

At this writing, there are several million people around the world who have access to a synthetic world. These worlds, known technically as MMORPGs (massively multiplayer online roleplaying games), are environments that allow people to undertake various tasks, hunting, socializing, exploring, producing and consuming goods and generally leading a more or less full, rich and detailed life. They have their roots in the text-based multi-user domains that date back to Richard Bartle's MUD1, deployed in 1978. Now, as then, many 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 (Castronova, 2001a). Since that time, at least 10 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, labour 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.

However, further thought suggests that virtual economies may be anything *but* normal.[1] 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 were demanded, at no cost to itself? If those two acts were possible, then a policy of government price control would be feasible. 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 preceding 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 multiplayer 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?

This paper attempts to address these questions, beginning in Section II with a simple rational choice model for determining the demand for game time. Given the unique features of gaming as a consumer good, Section III assesses possible market structures in the games market: will one game eventually dominate the world? Section IV considers the macroeconomic implications of large-scale expansion in the gaming phenomenon, especially for GDP and the tax base. Section V lays out some of the policy issues that widespread gaming will raise. Section VI 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. 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 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.[2]

Absent of 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. Following 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.

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.[3] 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 jigsaw puzzles. A puzzle with 900 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 jigsaw 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 behaviour that asserts a universal conflict between our ends and our means. The essence of behaviour, 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 behaviour 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 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 first assume that emotional well-being is always one goal of human behaviour. People do things that make them feel happier. Second, 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) $S = aR - b(C - W)^2$

where W 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.[4] When it comes to games and puzzles, the choice involves the amount of time spent 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) $U(H_A, H_B) = S_A\ln(H_A) + S_B\ln(H_B)$

where H_A and H_B are hours of play in the two games. If total time available is denoted by 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 act 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 gameplay: repeatedly playing the same game gets 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.[5]

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 gameplay.[6] 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) $G = Y - p_AH_A - p_BH_B$

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

4) $U(H_A, H_B, Y - p_AH_A - p_BH_B) = S_A\ln(H_A) + S_B\ln(H_B) + \gamma\ln(Y - p_AH_A - p_BH_B)$

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 that market, specifically its interaction with real world labour 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, the paper documented the fact that people can make real money by selling the digital

items that they produce while playing, and the average wages of game players were somewhat below average for Earth workers of similar education levels. 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 now to devote time to work time, denoted by L, and time in a single game, denoted by H. In other words, now we abandon the separation of gaming time into several games, and focus just on gaming versus work. Let game time have a price p, and produce satisfaction S, and let work be compensated at the wage rate w. Then:

5) $Y = wL$

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

6) $Z = T - H - L$

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

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

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 related to 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 by C, that are presented by your boss, your co-workers and your competitors, and where overcoming those challenges garners you rewards, denoted by R, in the form of wages, perks, fringe benefits and assorted entertainments involving the office copy machine.[7] 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 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 satisfactions of Earth work are directly compared to the emotional satisfactions of game time. These Earth/ game substitutions involve real economic transfers. It has already been shown (Castronova, 2001a) that labour 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.

III. 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.

These industries produce three items of interest, namely, connections, interface and content. Developments in connections include the internet and, increasingly, wireless communications. Development of interfaces includes voice command, head-up displays and body motion detection (computer-controlling gloves, gaze readers). Developments in content include the supply side of the market for games, where annual revenues have grown beyond Hollywood box office revenues.[8] 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 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.[9] 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 labour-intensive activities. One could perhaps increase production of content by allowing other producers (say, by opening game code to the public), but continued control of the world being created would be problematic.[10] 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. Sometimes the only way to reduce congestion is to add content, but this, again, is labour 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 gameplay around the time-intensive development of avatar capital. A player starts the game with a weak avatar, but gameplay 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, defenceless 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. On the other hand, two

games (Ultima Online and Dark Ages of Camelot) now offer methods to effectively start out ahead: in Ultima, you can directly buy your levels; in Camelot, you can start a new avatar at level 20 if you have already gotten one to level 50. These strategies help companies discourage the buying and selling of avatars outside the game, perhaps at a cost to the atmosphere within the world. 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.[11]

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.

This analysis allows a tentative answer to the first question of the study: in the medium-term future, the online multiplayer gaming market will probably consist of a number of large, densely populated worlds, with varying degrees of portability between them. The worlds will generate large revenue streams and will occupy many hours of human time, some of it considered play, some of it considered work. The hours that people devote to games will result in the accumulation of stocks of digital capital goods. These objects will have considerable economic value. Given the expected growth in connectivity, interface technologies and content, there is reason to believe that this 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?

IV. 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. It is important to recognize from the start that the mere fact that Earth economies may suffer as people spend more time in

cyberspace, does not imply that humanity is worse off. The fact that labour 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. A proper accounting would show, in fact, that the actual production of well-being 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 (e.g. 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 labour 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.

V. Constitution and Governance

Indeed, the most salient current policy issues both within and outside of games involve issues of governance. In the United States, there have been judicial rulings indicating that Earth courts have no jurisdiction over events that occur online (Kaplan, 2001). An argument (recently successful 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. Legal scholars have long recognized the formation of law within virtual worlds (Mnookin, 2001; Dibbell, 1999). Now they have begun to address broader issues involving the intersection of virtual world activities and Earth law (Lastowka and Hunter, 2004). 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 (e.g. 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. "Foreign trade" refers to the common practice of selling in-game items for real money in out-of-game markets such as eBay. This trade is simple to conduct and hard to detect.

It also 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 gameplay, and damages the entire gaming environment; the situation is a commons tragedy, where the self-interested

trading behaviour 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. Users of early text-based worlds certainly conceived of themselves in this way (Mnookin, 2001; Dibbell, 1999). 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. 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. 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 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.

VI. 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 behaviours 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 in 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.

Notes:

[1] The first people to compare and contrast virtual and real economies include John Beezer and Zachary Booth Simpson, both of whom have unpublished analyses of game economies on the internet. Richard Bartle's vast experience with game economies is now available in his textbook (Bartle, 2003), which is strongly recommended for anyone thinking about actually building one of these places.

[2] 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 Neumman - 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 (for example, 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 *SimCity* and its spin-offs. I suspect that many of my colleagues in the social sciences have played *SimCity*, 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 roleplaying this job was worthy of further investigation. Nor did it occur to me that this behaviour, in itself, might reveal something fundamental about the economic and social behaviour 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*.

[3] Police and courts in Korea now regularly get involved in cases of theft of virtual world items. In the United States a mock trial held before US Circuit Court Judge Philip M. Pro on July 30, 2003 (of counsel: Richard Salgado, Department of Justice and Jennifer Granick, Stanford Law) determined that the loss of US$5,000 in digital game items would constitute a real loss, sufficient to trigger a potential felony prosecution under US Code Title 18, Section 1030.

[4] 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 be the things that people pursue, the converse does not necessarily hold. People may well pursue things that do not make them feel good. Certainly, anyone who works in the therapeutic industries can confirm that many people devote significant amounts of energy to behaviours that they quite consciously know will not make them happy. In theoretical terms, this idea can be captured if we 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. Obsession with work and money is a well-known counterexample in economics (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 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.

[5] 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. Profit-maximization entails equating marginal revenues to marginal costs, so, if the marginal cost of challenge is positive, profit-seeking will force game companies to choose a challenge level where marginal revenues are also positive. At this optimum point, the marginal contribution of challenge to emotional satisfaction must also be positive. And this will only happen at values of C that are below W. Thus, in competitive markets for puzzles, all puzzles will be at least a little bit too easy.

[6] 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.

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

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

[9] 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 equalled or surpassed the entire development cost.

[10] This is the strategy of Project Entropia (Mindark) and Neverwinter Nights (Bioware).

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

References

Au, Wagner James (2002), "Playing Games With Free Speech," Technology and Business, salon.com, May 6.

Bartle, Richard. *Designing Virtual Worlds*. Indianapolis: New Riders, 2003.

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.

Dibbell, Julian. *My Tiny Life: Crime and Passion in a Virtual World.* New York: Henry Holt, 1999.

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.

Lastowka, F. Greg and Dan Hunter (2004), "The Laws of Virtual Worlds," forthcoming in *California Law Review*.

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.

Mnookin, Jennifer L. "Virtual(ly) Law: The Emergence of Law in an On-Line Community." In Peter Ludlow (ed.) *Crypto Anarchy,*

Cyberstates, and Pirate Utopias. Cambridge: MIT Press, 2001: 245-302.

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.

[To the top of the page]