Analysis of real markets in a synthetic world

András Margitay-Becht (Saint Mary's College of California, USA)

Abstract

In theoretical economics we frequently make use of very strong axioms in order to simplify our models. We recognize humans are not perfectly rational, but the rationality assumption is so useful in predicting behavior, we employ it all the time. Similarly, while acknowledging that information is neither free nor freely available, many times we make use of the assumption that every agent in our economy has perfect information. And while we understand that many of our results rely too heavily on them, we have to use them, as the real world is too complex to analyze otherwise.

Over the past year, I have been working on an alternative approach, something that I have already used successfully in the classroom. The past two decades has seen the rise of immersive virtual realities, which can be understood as a kind of telecommunication-based society with voluntary membership. Some of these virtual realities became expansive enough to form actual worlds, where real world humans interact with each other, plan, produce, buy and sell goods – live, at least part time. Some of these realities have populations of 10, even 20 million people.

Since the summer of 2013, I have been collecting daily transaction data of all trades being offered in a virtual community of over 2 million Europeans in a set of roughly 250 synthetic worlds. This vast amount of data allows for intertemporal, cross-cultural, spatial and other comparisons that real-world market data does not make possible. In these synthetic worlds we do not need to account for the actions of overbearing monopolies and oligopolies with powerful marketing campaigns, we can observe economics work among real people. In this paper I will show that this data can serve as relevant foundation to expanding economic understanding.

Introduction

This paper deals with preliminary research, primarily aimed at determining the relevance and usefulness of following this course of inquiry. My primary aim at this stage is to show that real transaction data from a synthetic setting does provide insights that are applicable to the real world.

In the paper I first introduce the concept of the synthetic worlds, and their relationship with real-world economies and economics itself. Then I will describe the proposed research itself, and the technical details pertaining to the gathering and processing of the data. After demonstrating some of the results from this preliminary inquiry, I conclude with the next steps in the course of this research.

What are synthetic worlds?

Synthetic worlds are computer-based simulated environments. The term is coined by Edward Castronova (2005) to describe all virtual worlds, emphasizing that they are real, just not organic, rather than synthetic. These modeled worlds have their own set of rules, landscapes, and participants.

A subset of synthetic worlds are the so called Massively Multiplayer On-Line Role Playing Games or MMORPGs for short. These games are designed to be immersive, frequently three-dimensional virtual environments, which are engaged by the users through the use of avatars: computer models that they themselves control. They are called “massively multiplayer”, as they serve as meeting grounds for thousands or tens of thousands of players at a given time. This interactivity expands the impact of these games from being purely a medium of entertainment: they become an environment for social interactions that end up becoming an important theater for personal development (Yee, 2006).

Synthetic worlds and economics

Synthetic world and economics has been intertwined in a number of ways over the past decade. One avenue of interaction has real-world businesses use virtual settings to sell their products, have meetings, or hold promotional events. Second Life is especially popular for this regard, as this virtual world is nothing but a virtual sandbox. Players enter with their avatar, and navigate the landscape owned and designed by their fellow players. Linden Labs, the company operating this virtual world, makes their money from selling the virtual land, and all "content" is created by the users. This setup makes it easy for real-world entities to present themselves in this environment, and many decide to do so. On May 22, 2006, for example, 20th Century Fox held a premiere for X-Men: The Last Stand in Second Life (James, 2006). Cisco, American Apparel, Dell, Disney, Mazda, Sun, Toyota and Wells Fargo are just some of the companies that have at one point operated within Second Life. Educational institutions also took interest. Harvard Law School set up a college campus as early as in 2006 (Wong, 2006). My own home institution, Saint Mary's College, has owned and operated a Second Life island for educational purposes for over half a decade now, where we hold meetings and at times hybrid classes.

While synthetic worlds impacted real-world businesses, real-world business processes and economics has also impacted virtual worlds. Many virtual worlds offer a marketplace for their users to exchange goods with each other – and they have found that albeit the goods exchanged are virtual, the economic processes governing these exchanges are all too real. In 2007, CCP Games hired Dr. Eyjólfur Guðmundsson to study the complex economy of EVE Online virtual world (Hillis, 2007) – a job he held until July 31 of this year, when he became the rector of University of Akureyri. In 2012, Valve hired Yanis Varoufakis (Griffiths, 2012), though his public presence at the company was somewhat short-lived.

On the other side of things, the users of the virtual worlds themselves apply economic theory and finance to understand (and try to predict) the way markets work in their virtual society.

Proposed research

The initial motivation behind my research was the above two paths of interaction between synthetic and real. If real-world entities can use the virtual medium to spread their message and to teach, and if the virtual worlds can use real-world economics to further their efficiency, maybe it is possible to use the *data* from virtual realities to conduct real-world applicable big-data economic research. The medium of choice is World of Warcraft (WoW), Blizzard Entertainment's flagship product that has dominated the MMORPG markets in the West over the past 7 years. This fantasy-themed game is based in a cartoon-like virtual environment, where players belong to one of two warring factions (the Alliance and the Horde), and engage either computer-controlled enemies or each other. The content players have at their disposal is rich and varied. They explore a large virtual world (estimated at roughly 130-150 square miles, based on Stoutfoot (2007) and extrapolation), fulfill quests given by friendly characters, and essentially living through, hundreds of greater and thousands of lesser storylines. Aside from this interactive narrative, participants can also tackle greater challenges that require groups of 5-10-20 or even 40 like-minded players to band together and face. The game also provides content for those who like to collect things, enjoy turn-based strategy games or exploration.

When creating their avatars, the players select their race – from something as similar as a human to as outlandish as a humanoid cow (in World of Warcraft terminology, a Tauren), and their "class" – a set of abilities that define how said avatar interacts with the environment. On top of that, each avatar can choose up to two professions or trade skills, which enable them to gather or craft materials and goods for use of themselves – or for sale to others. To facilitate the sales, Blizzard created an auction house (AH) system: players can place goods for sale on their server's auction house, and other players can browse these listings and purchase the goods. It works very similarly to how E-Bay handles transactions: people can bid on goods, and when the time is up, the highest bidder receives the product. Or you can choose to pay the buyout price and receive the product immediately. Our economic inquiries are going to be focused on these auctions.

Why World of Warcraft?

There are a number of reasons why using World of Warcraft as our data source makes sense. The biggest and most obvious one is that Blizzard Entertainment adopted an innovative policy when it came to user created content: they provided an application programming interface for enterprising users that makes it possible to access, gather and display data. On the one hand, this makes it extremely easy to do the actual data gathering: writing a simple script one can connect to Blizzard's servers, and retrieve the auction house data in a matter of minutes, even without logging in to the game itself.

This cavalier attitude to user-created programs, however, has other important implications, too. One of the big, and frequently unrealistic assumptions economics makes when discussing markets is the assumption of "perfect information": every agent knows everything. They know the products, the current prices, historical prices, processes, everything. Because of the user-created programs, this assumption in World of Warcraft is not unrealistic.

- The users know the products. Every product sold is designed by Blizzard itself; has a unique type ID, and every instance of a given product is identical. When buying Iron Ore, the World of Warcraft consumer does not need to worry about purity, quality, joules/kg; an Iron Ore is an Iron Ore
- The users are aware of the price for every single seller. They are able to list every single auction house listing for a given product, and can order that list by price. Some of the player-created programs, called "add-ons" or addons, are able to pull the entire database, and even store it for later use. These programs not only show their user the current value of a given product, but also show it in the context of historical data: how good of a deal they are getting right now.
- The users are also familiar with the processes of creating the items. If they are not, there are many third-party web-sites available where they can easily get this information. Wowhead.com, for example, lists every single item, all the forms of accessing it, and all of their potential uses as well.

Because of this wide availability of information, the assumption of perfect (or near-perfect) information is going to be much closer to reality, than it is in the real world.

World of Warcraft is also useful as the synthetic world of choice because it is not a casual game. The average gamer spends 5 hours a day playing, is 30 years of age, and many of them are females. While subscriber numbers have been declining over the past few years, WoW still has more than 6 million active subscribers. In my research I focus on the 265 European servers (housing roughly 2 million players), as this adds a new dimension (culture) as a possible line of inquiry.

The Data Gathering and management considerations

To gather the data, I use a custom-written java program. This connects to Blizzard's servers, from where it first downloads the timestamp of the most recent data dump. If this timestamp is later than the last datapull, the program downloads the updated auction house data. Currently I store this downloaded data in its raw form, as a JSON format text file. In order to be able to place many high-speed inquiries, it would be convenient if this data could be loaded into a relational database; however due to the large amount of data this would only be realistically feasible if the database was housed on high-speed SSD drives.

Because of this technical limitation, the queries run on the database use my own "text-based" database query system, which works, but cannot be considered high-speed. One of the goals of this preliminary research is to determine if it is worthwhile to spend thousands of dollars on a high-speed data architecture.

<u>Few sample results</u>

There are many potential research directions to pursue with this relatively large dataset. Even with the constraints of hardware and bandwidth, my database currently holds roughly 13 billion datapoints (which could be loaded into a relational database in about 50-55 days at my hard-drive speeds). To help decide whether there is any point, I ran a few sample queries on product 72092, "Ghost Iron Ore" on some servers. I chose this product as it is one of the most commonly used raw materials in the current World of Warcraft marketplace.

Market size

The drop in subscriber numbers opened up an avenue of research for economists. Since some serves lost a large chunk of their populations, it became hard to trade for people. To alleviate this issue, Blizzard is slowly merging the auction houses of certain low population realms. This, of course, allows us to observe what happens when market size increases. In the examples below, you will see the Ghost Iron Ore market data from 3 of the 265 servers in Europe, from December 2013 to the end of March 2014. These dates were chosen as at the beginning of January the auction house of Tyrande became connected to Colinas Pardas, and at the end of February Los Errantes was connected to them as well.

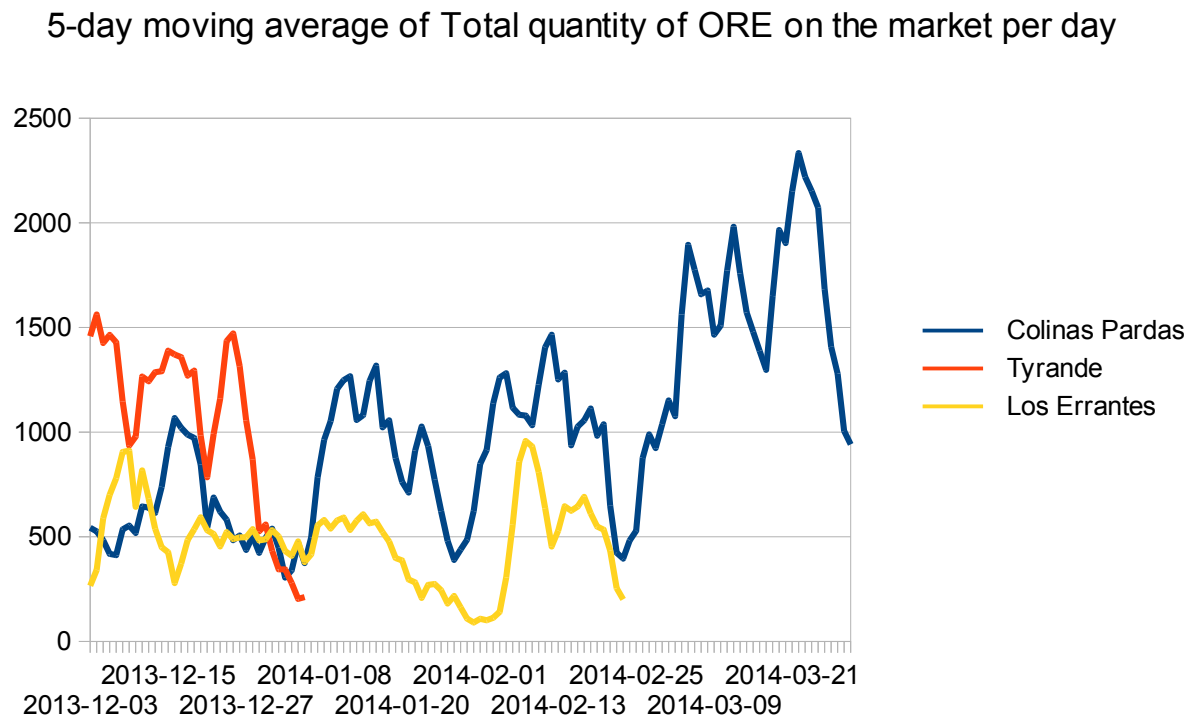

In the above diagram we can see that with the merger of the servers, the amount of available ore has increased – as expected. More buyers, more sellers, the blue line keeps going higher as it swallows up orange and yellow. But what happens to price?

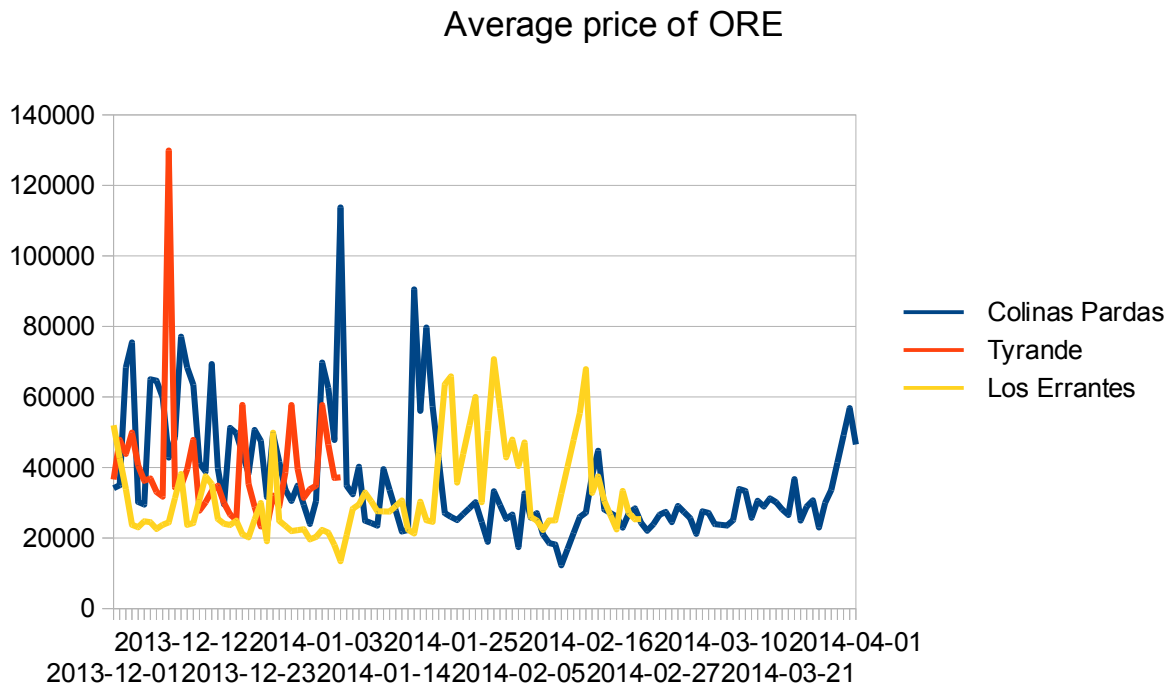

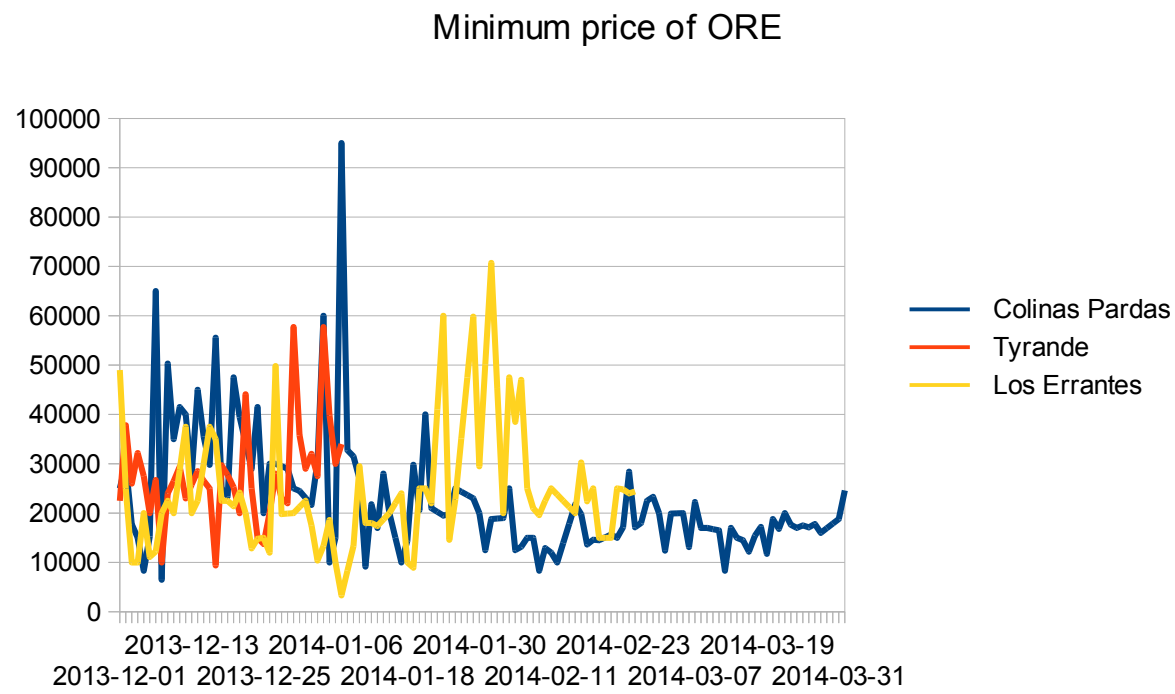

As we can see, both the average and minimum price of the ore decreases. This is also in line with our expectations: the greater, more competitive market drives down the prices. This also means that the competition among the suppliers seems to be greater than the competition among the buyers. Parallel to this, we can also observe a drastic decrease in price volatility.

Quantity supplied versus price

Since we have the price and quantity data for every auction, we can look at the relationship between quantity of ore offered, and the average or minimum price of the ore on the market.

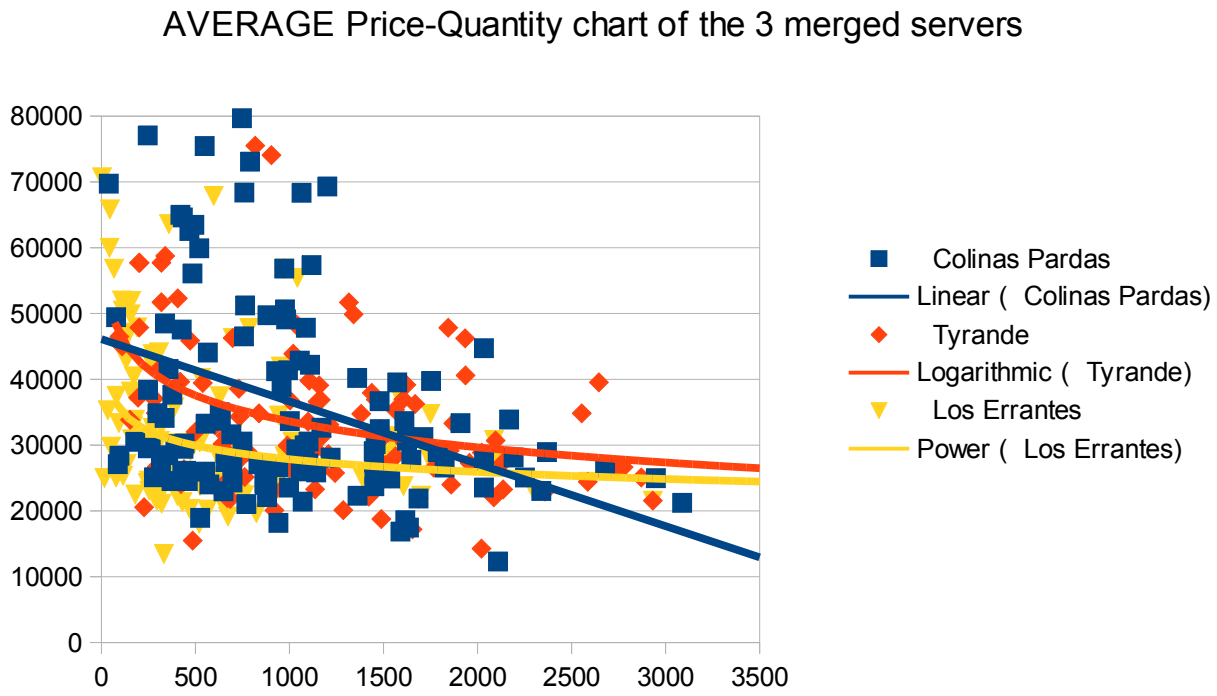

It is clear to see that there seems to be a weak negative correlation between average price and quantity offered.

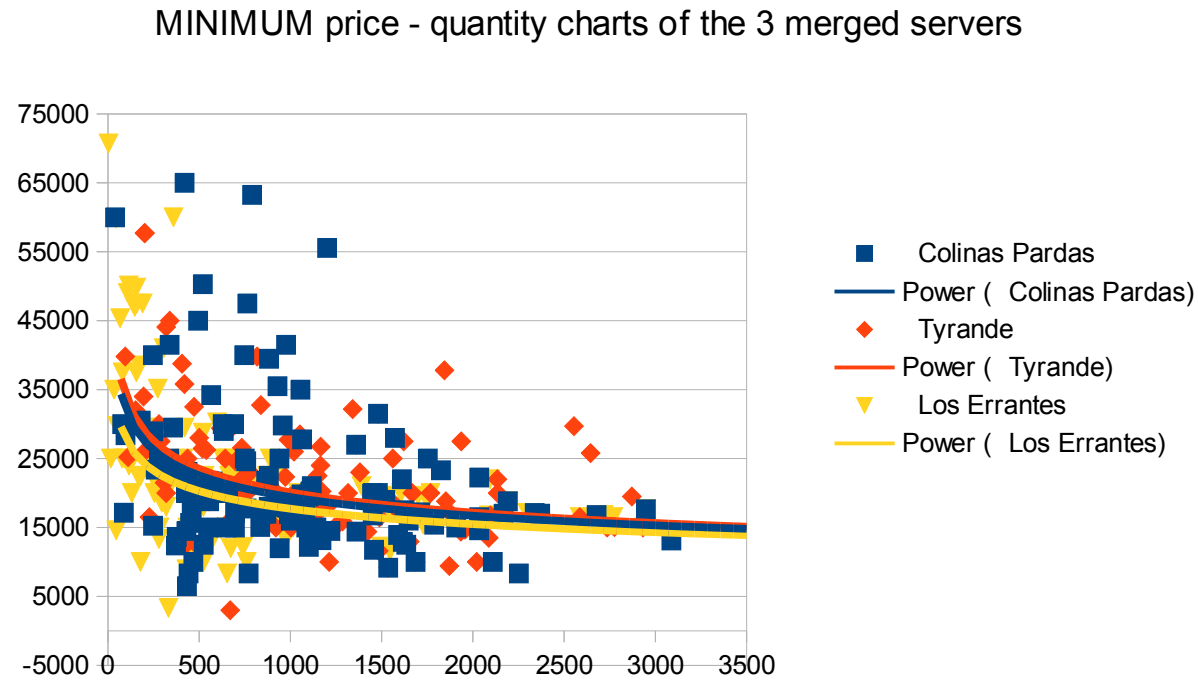

The correlation between minimum price and quantity of ore offered seems to be even weaker, but present. Economic theory would expect a much stronger correlation between these two measures.

Conclusion and further research

The above few sample diagrams show, that virtual economy data can hold interest to us regarding the real world. The first example dealing with market size increase is a good illustration to the transitioning from a closed-economy to an open-economy setting. The results are encouraging. On the other hand, we seem to have found some curious data about quantity supplied and price. What could cause this weak negative connection?

Using the available data, we can run nearly endless simulations to explore much more complex questions. Among others, we can:

- Calculate Hirshman-Herfindahl indexes for all sellers in a market, and then compare the behaviors of markets based on this concentration figure. How strong is the relationship between market concentration and lower efficiency?
- Calculate and compare average and minimum prices of certain goods across all servers
- Calculate a price level for a given server – or across all servers
- Calculate inflation in a segment or total
- Determine the impact of speculation on price (price vs. HH)
- Get a broad understanding of the connection between market size and market efficiency. Will larger markets in general lead to more effective markets (as in our preliminary finding), or will it just lead to more powerful oligarchs?
- Determine whether cultural values appear in market prices. Are some nations in Europe running more efficient markets than others? We have separate data for French, German, Italian, Russian and Spanish servers, in addition to the "everything else" English-speaking servers

Pursuing these inquiries can deepen our understanding of not only economics, but how these disparate aspects are related to each other.

Bibliography

Castronova, E. (2005). Synthetic Worlds: The Business and Culture of Online Games. Chicago. University of Chicago Press.

Griffiths, D.N. (2012). The Value of Fun: Valve Software Appoints In-House Economist. Reuters. http://www.forbes.com/sites/danielnyegriffiths/2012/06/15/valve-appoints-in-house-economist/

Hillis, S. (2007). Virtual world hires real economist. Reuters. http://www.reuters.com/article/2007/08/16/us-videogames-economist-life-idUSN0925619220070816

James, W (2006). The Uncanny X-men come to second life. New World Notes. http://nwn.blogs.com/nwn/2006/05/the_uncanny_xme.html

Stoutfoot, T. (2007). How big is Azeroth? (Pseudonym) http://tobolds.blogspot.com/2007/01/how-big-is-azeroth.html

Wong, G. (2006). Educators explore 'Second Life' online. CNN. http://www.cnn.com/2006/TECH/11/13/second.life.university/index.html

Yee, N. (2006). The Psychology of MMORPGs: Emotional Investment, Motivations, Relationship Formation, and Problematic Usage. In R. Schroeder & A. Axelsson (Eds.), Avatars at Work and Play: Collaboration and Interaction in Shared Virtual Environments (187-207). London. Springer-Verlag.