

A Test of the Law of Demand in a Virtual World: Exploring the Petri Dish Approach to Social Science

EDWARD CASTRONOVA

CESIFO WORKING PAPER NO. 2355
CATEGORY 10: EMPIRICAL AND THEORETICAL METHODS
JULY 2008

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

A Test of the Law of Demand in a Virtual World: Exploring the Petri Dish Approach to Social Science

Abstract

We report results of an experiment on prices and demand in a fantasy-based virtual world. A virtual world is a persistent, synthetic, online environment that can be accessed by many users at the same time. Because most virtual worlds are built around a fantasy theme, complete with magic, monsters, and treasure, there is considerable skepticism that human behavior in such environments is in any way “normal”. Our world, “Arden”, was designed to test whether players in a typical fantasy environment were economically “normal.” Specifically, we tested whether fantasy gamers conform to the Law of Demand, which states that increasing the price of a good, all else equal, will reduce the quantity demanded. We created two exactly equivalent worlds, and randomly assigned players to one or the other. The only difference in the two worlds was that the price of a single good, a health potion, was twice as high in the experimental world than in the control. We allowed players ($N = 43$) to enter and play the environment for a month. We found that players in the experimental condition purchased 43.1 percent fewer of the potions, implying a demand elasticity of -0.431 . This finding is well within the range one expects for normal economic agents. We take this as evidence that the Law of Demand holds in fantasy environments, which suggests in turn that fantasy gamers may well be economically normal. If so, it may be worthwhile to conduct controlled economic and social experiments in virtual worlds at greater scales of both population (thousands of users) and time (many months).

JEL Code: C99.

Keywords: virtual worlds, experimental methods.

Edward Castronova
Department of Telecommunications
Indiana University
1229 E. 7th Street
USA – Bloomington, IN 47405
castro@indiana.edu

*Author list:**
Mark W. Bell, Marc Carlton, Robert
Cornell, James J. Cummings, Will Emigh,
Matthew Falk, Michael Fatten, Paul
LaFourrest, Nathan Mishler, Justin Reynard,
Sarah Robbins, Travis Ross, Will Ryan, Rory
Starks

12 July 2008

This research was made possible by a grant from the MacArthur Foundation. Bioware, Inc., makers of the *Neverwinter Nights* videogame and the *Aurora Toolkit*, provided free software for use in the experiment. *The author list includes all members of the team that built and tested the experimental environment. Castronova was the principle investigator; the other contributors are listed in alphabetical order.

1. Introduction: Virtual Worlds As Petri Dishes

Since 1997, there has been a rapid increase in deployment and use of virtual worlds: online computer-generated environments that can be accessed by thousands or even millions of people at the same time (Castronova 2005). Each user explores the virtual world using an *avatar*, a human or human-like character that walks, talks, fights, and hugs, much as a real human does. While inhabiting her avatar, the user can chat with other people, play games, go on quests, or harvest virtual goods. The range of activities is determined by the designers and is typically vast since, unlike real-life people, an avatar can fly, turn into a fish, or become purple, just by changing the underlying computer code.

With this kind of design flexibility, there are many possible applications for a virtual world. While most are devoted to entertainment, some worlds are being used for serious purposes: to train first responders, teach business lessons, or treat minor mental disorders.

In this paper, we consider a serious research application of virtual worlds: their use for controlled experiments. Several unique features argue for this kind of use. First, virtual worlds allow controlled experimentation at the level of an entire world. If desired, designers can fix the code so that two research environments are exactly the same, down to every leaf on every tree. Second, this technology allows truly vast research environments. If desired, a research team could create a world that covers hundreds of millions of square miles. Third, virtual worlds allow huge numbers of research subjects. Current commercial virtual worlds commonly have many millions of players. Fourth, virtual worlds allow a long time scale for research. Again, current commercial virtual worlds typically retain users for many years. One world, *Ultima Online*, has been in continuous existence since 1997. Finally, the people who engage with one another inside virtual worlds seem to constitute a genuine society. Casual observation reveals that, even at modest size, significant social relationships seem to develop. It appears, again to the casual observer, that friendships and reputations are forged and broken. Some individuals seem to acquire some kind of interpersonal or political power, while others are deemed “weak.” Information networks seem to be active. Perhaps most surprising, virtual worlds seem to develop internal markets, in which players trade virtual items with one another at what appear to be stable and robust prices.

More careful observation of virtual worlds tends to confirm these suspicions. In a recent study, Chesney, Chuah and Hoffman (2007) conducted a series of classic experiments from experimental economics within the environment of *Second Life*. They found that, for the most part, the usual environment of a small-scale social science research lab can be replicated in a virtual environment. Almost all aspects of the usual experiment can be recreated virtually. The one exception to this finding was in the area of physical signals and cues, which of course are not transmitted by avatars (at least with current technology).

Results such as these raise the possibility that virtual experimental environments might be possible on an even greater scale. In an earlier paper, an effort was made to estimate the real-world productivity of characters in virtual worlds (Castronova, 2001). Real-world transactions of virtual currency for US dollars yielded a shadow price for the virtual currency. This in turn was used to translate virtual values into real values. By this kind of accounting, it appears that the gross level of economic transactions in virtual worlds has already grown into the billions of US dollars per annum (Castronova, 2007).

Similarly, anthropologists have said that they observe genuine cultural behavior in virtual worlds, as well as between them and the real world (Malaby, 2006); sociologists see genuine group formation (Taylor and Jakobson, 2003); and legal scholars see the outlines of real-world law (Hunter and Lastowka, 2004).

With the exception of Chesney et al., all of these conclusions are based on direct observation. The method is that a reputable and careful scholar enters a virtual world, collects data (formally or informally) and then reports that the data seem to follow the same patterns in the virtual world as in the real world. No particular theory of behavior is tested. With the Chesney et al. study, the focus was on replicating the typical economics experiment lab, not society on a grand scale.

This leaves the real-world legitimacy of the large-scale behavior we see in virtual worlds in a state of doubt. Is the economy of such a game-world a real economy? It remains possible that, for most players, their participation in the virtual world is “just a game” and therefore any large-scale patterns are not analogs to real-world counterparts. Indeed, the stated object of participation in most virtual worlds is precisely that, to play a game, to behave in a way that is new, to experience things that are not possible in the real world. The mere fact that people trade things at certain prices, or inhabit this or that social role, need not imply that people are behaving according to our accepted economic and social theories. Perhaps when a man is an elf or an orc, he decides to buy high and sell low instead of the more sensible obverse pattern. Perhaps, when a woman play-acts a ship captain, she defers all decisions on the ship’s course to the bilge crew, instead of making them herself. And why not? In a fantasy environment, it does not matter whether a person has enough money to “live” on, or whether a ship gets to a certain port “on time”.

The previously cited observers of virtual worlds would counter with broad evidence that even in play-acting environments, human behavioral theories continue to hold water. The Stanford Prison Experiment and its follow-ups strongly suggest that people who are consciously play-acting a role nonetheless behave as if the role was real (Zimbardo, 2007). The fact that players of virtual worlds have committed murder and suicide over events within them certainly

suggests that these events matter to those involved (though obviously not healthily so), despite the fact that the worlds are fanciful by design.¹

This evidence is all indirect, however, and thus the question of whether virtual world behavior is “true” human behavior is open. Researchers in communications have tackled this question directly, however. Byron Reeves and Clifford Nass summarize a long series of brain experiments showing that people seem to treat media as real, apparently because of the simple fact that the brain evolved before media existed (Nass and Reeves, 1996). More recently, Nick Yee and Jeremy Bailenson (2007) have coined the term “Proteus Effect” to describe the phenomenon of a person treating his or her virtual body as if it were a real body. In their studies, for example, people given a taller avatar in a virtual world act, in that world, more confidently. That is exactly the correlation we observe in the real world, namely, that height predicts social confidence. Yee and Bailenson demonstrate that this feature of the human evolutionary make-up translates directly into the virtual environment.

Yee and Bailenson’s notion of a Proteus Effect can be studied in other contexts. Here, we ask whether there is a Proteus Effect involving economic calculus. One can easily imagine (and the authors expect to see) a stream of research showing where human behavior generalizes across the virtual membrane and where it does not.

Determining where virtual world behavior mimics real world behavior is quite important for methodological reasons. If virtual world behavior can be treated as a model of human behavior in general, this would allow a fresh approach to empirical social science. In the natural sciences, empirical work is based largely on a foundation of controlled experimentation. In the social sciences, controlled experiments can generally only be conducted on the most limited scale – a small number of people, over a small time period, for minor rewards. Yet the most pressing problems often involve large numbers of people in persistent environments. Researchers who have wrestled with this problem in the past have generally assumed that it would be impossible to conduct experiments that are truly macro-level and persistent. In the words of economist Robert Lucas, “The problem involved in convincing a collection of experimental subjects that they are in an infinite-horizon environment seems to me insurmountable” (1986, p. S421). Such a problem may have been insurmountable in 1986, but virtual worlds provide a potential solution.

With virtual worlds, the experimental environment can be made to persist quite literally forever. Already, in commercial versions such as Blizzard Inc.’s *World of Warcraft*, we have millions of people, over the course of years, pursuing rewards of such value that they can even motivate some people to kill themselves or others. In such an environment, it would be possible to conduct controlled

¹ “Ill Hudson Man Took His Own Life After Long Hours On Web,” Milwaukee Journal Sentinel, March 31, 2002. “‘Game Theft’ Led to Fatal Attack,” BBC News, March 31, 2005.

experiments at the macro level. Natural experiments that have already occurred seem to demonstrate the feasibility of this approach (Castronova, 2006).²

Such a method overcomes the difficulties faced by social researchers in establishing large-scale causation. In a typical regression analysis, one uses independent variables to attempt to isolate the direct relationship between the variable of interest and the dependent variable. However, even a lengthy list of independent variables is insufficient to completely isolate the effect of interest. Moreover, most regression analyses must make use of quite a bit more artillery – multiple equation systems, adjustment for unobserved variables, different functional forms, and so on. But none of this, of course, is sufficient to establish causation; it merely isolates the direct correlation between a specific independent variable and the dependent variable. This isolated relationship is then assessed for whether or not it is consistent with a casual theory. Contrast this complex and indirect method of inference with the far simpler method of controlled experimentation, on which the natural sciences are based. Place two equivalent Petri dishes on the counter. Insert the same yeast in each one. Expose one of the dishes to fire. Observe that the no yeast grows in the burned dish. Conclude: “Fire causes yeast to die.” Causation is directly identified, in a method that is simple, replicable, and persuasive. This is the method that social science might be able to use, in place of regressions and classical statistical inference, if virtual worlds are found to induce behavior that generalizes to the real world.

Testing for general behavioral theory in virtual worlds is thus a necessary first step in exploring whether virtual worlds can serve as social science Petri dishes. The objective of the experiment reported in this paper was to conduct such a general behavioral test.

Our test focused on the Law of Demand, which holds that as the price of a good rises, all else equal, the quantity demanded of it will fall. There is perhaps no principle of human behavior more universally accepted in all of economic theory. We know of no introductory textbook that does not state the Law of Demand in its earliest, and foundational, chapters. Thus a test of the Law of Demand is indeed a stern one for virtual worlds. If the Law of Demand does not hold in a virtual world, we may as well conclude that virtual worlds generate no economic behavior of general interest.

Our test was simple: We created two virtual worlds that were exactly the same in all particulars save one: in one version of the world, the price of a single good was twice as high as in the other. We then assigned subjects randomly to the two worlds and observed how many items were purchased. The Law of Demand requires that, since, indeed, all else was absolutely equal in this experiment, the higher-priced good should have been purchased less frequently. But since our virtual world had monsters, treasure, and magic in it, and since the players took

² An outbreak of “plague” in one virtual world was subjected to serious scrutiny by epidemiologists. See Balicer (2007) and Lofgren and Fefferman (2007).

on the personae of elves, hobbits, and dwarves, perhaps there was no reason for anyone to pay attention to economic rationality. If not, the Law of Demand would be a clear casualty. Our experiment was designed to test whether this Law survived the translation from the real world into an environment of dungeons and dragons.

In the next section, we will describe the environment we created. Section 3 describes the experimental protocol. Section 4 presents results, and Section 5 offers a concluding discussion.

2. The Environment: Arden

The virtual world supporting this experiment was named *Arden: The World of William Shakespeare*. Arden was built using the Aurora Toolkit, a virtual-world building software created by Bioware Inc. in support of its fantasy game *Neverwinter Nights*. The Aurora Toolkit allows the game designer to create a three-dimensional fantasy environment and populate it with non-player characters (NPCs), monsters, treasures, and quests. This environment can be hosted on a server and opened to traffic from the internet. People can then log in to the environment and choose a character (dwarven warrior, elven priestess, human thief, etc.). Using this character, they can then interact with whatever the play environment contains.

We created a world of 41 three-dimensional “zones”, or play areas, 46 unique NPCs, 23 unique monsters, 25 special items, 111 conversation files (programs to handle interactions between a player and an NPC), and 204 custom scripts (programs to execute game functions). In addition, we used hundreds of the generic monsters, items, and objects available in the toolkit.³

As an illustration of the kind of work involved, Figure 1 shows the code used to generate money in the game. This code would execute whenever a player opened an object labeled “Dropped Loot”. On each such object was a local variable set to “High”, “Medium”, “Low, or “Wild”, and depending on this variable, a certain number of gold pieces would be created within the object when it was opened. For example, a “low” item would generate “d12 + 4” gold pieces, meaning, a random integer between 1 and 12, inclusive, plus 4. Players in the world could walk their characters to the location of a Dropped Loot item and click on it. At this point the code would create gold pieces in the item, which the player could then take. The Dropped Loot item would then disappear for a time, reappearing later for a new player to come by and inspect. This in fact is how players obtained purchasing power: by harvesting gold pieces from the Dropped Loot items that appeared and disappeared from time to time. The Dropped Loot

³ A “Ring of Protection +1” is a generic item. Our “Hamlet’s Ring” offered a magical effect of our own design, and was therefore a unique item. Similarly, a “Bear” is a generic monster, while our “Escaped Bear” had its own AI routines and was therefore a unique monster.

items were allocated so that every player would have a substantial, but not overwhelming, amount of purchasing power throughout the game.

```
void main()
{
    int nGold;

    int nGoldAmount = GetLocalInt(OBJECT_SELF, "nGoldAmount");
    switch(nGoldAmount)
    {
        case 1: //Low
            nGold = d12() + 4;
            break;

        case 2: //Medium
            nGold = d12(2) + 8;
            break;

        case 3: //High
            nGold = d12(4) + 16;
            break;

        case 4: //Wild
            nGold = d20(8);
            break;

        default:
            nGold = d8(2) + 2;
            break;
    }

    CreateItemOnObject("nw_it_gold001", OBJECT_SELF, nGold);

    PlaySound("it_coins");
}
```

Figure 1. Code For Generating Gold Pieces

The theme of our world was taken from William Shakespeare's *Richard III*. On entering the world, players were told (by a non-player character or NPC, named "Sergeant Bridgeford") that fires had broken out all over London and rioting had followed. They were then given a quest to help quell these riots, by order of King Richard III. Players then made their way across our virtual London, avoiding burning buildings as best they could while fending off attacks from mobs of thieves and looters. In the course of these attacks, they came across our Dropped Loot items and acquired the aforementioned gold pieces, the standard currency of the game. Upon reaching Newgate Prison, the epicenter of the riots, the players were told (again by an NPC) that they must go into the prison and kill

the rebel leader, Rydderch ap Rhys. In doing so, they discovered certain information about the King's top advisor, William Tyrrel, revealing him to be quite unsavory. They also discovered evidence that the troubles in the city were directly linked to an outbreak of restless spirits in the Tower of London. The players then made their way to the Tower and battled the undead spirits until, in a big showdown deep within the Tower's crypts, they learned the truth about King Richard: that he is a tyrant and a murderer (as depicted in the play), who uses Tyrrel as his foul henchmen. Tyrrel's worst act: the murder of two young boys. Their ghosts then appeared to the players and pleaded that the players rebel against this evil King. Tyrrel himself then appeared and demanded that the players support law and order and kill these troublesome ghosts. Each player then chose whether to fight Tyrrel or the ghost-boys, thus placing themselves either on the side of the Lancastrians or the Yorkists in the historical context of Shakespeare's play. This finished the game.

All told, a player might require 10 – 15 hours to complete this story; differences in hours played depend on player skills and tastes for exploration and pace. The environment included side quests and ancillary activities (social chatting with other players) that allowed players to vary the amount of time spent in the world. A player could enter the world at any time and exit it at any time; the world remained "switched on" around the clock, so that when the player returned, her character would be exactly where she left it. Players could experience the environment alone or as a group; a group could be formed in advance, by prior arrangement, or it could be random, as when several players happened to enter the same play zone at the same time. By design, a virtual world allows players to participate at whatever pace and in whatever way they wish, and this was the case in Arden.

As players made their way along, they encountered merchant NPCs, where they were able to use their cash to buy upgrades for their characters, such as stronger armor and sharper weapons. Each merchant offered hundreds of different items, at a wide variety of prices and features. The gold coin treasures were placed in the world so that a typical player going up to any merchant would always find many items that he could afford, as well as many that he could not afford. At the earliest levels, players would encounter enough Dropped Loot so as to acquire a minor endowment of 200 – 300 gold pieces fairly quickly. By the end of the game, a typical player might have 5,000 – 7,000 gold pieces. Of course, the items a player needs also grow more powerful, and more expensive, as the game progresses.

For example, among the most useful items in a game such as this are potions. By drinking a potion, a character could become invisible, gain increased armor, or learn the magical properties of an item. Health potions allow a character to restore her hitpoints, or HP. Hitpoints are a measure of the remaining stamina of a character during a battle. A character might start the fight with 20 HP. Each time the opponent makes a good blow, the character might lose 2, 5, or 10 of

these HPs, reflecting the vigor required to fend the blow off. If an attack takes the character below 0 HP, she is considered defeated. “Defeat” here means not death, but rather being teleported back to a starting point of the zone. The character now has to fight through all the zone’s monsters again in order to continue on with the game. Losing a fight is a setback in the player’s progress.

A health potion helps during fights because it instantly cures HP. However, how much help it is depends on the character’s level. A level 1 character might have 12 HP, meaning that a Potion of Cure Light Wounds, which heals up to 8 HP, can be quite useful. A level 10 character, however, might have more than 100 HP, and would typically be facing enemies that cause 15 or 20 HP of damage with each attack. For this character, a Potion of Cure Light Wounds has little use: the damage it cures is destroyed again in the time it takes to drink the potion. A Potion of Heal, which restores all HP, would be much better. Of course, Heal potions are many times more expensive than Cure Light Wounds potions. Thus, as a character levels from level 1 to level 10, her cash reserves increase, but so does the cost of operations.

In sum, the environment of Arden expressed most of the standards of fantasy virtual worlds. Players entered the world and navigated it as a kind of role-playing character; they received quests; they obtained treasure, magic items, powers, and spells; they encountered monsters and dungeons. Under all of this was a simple system by which players could obtain cash and buy items that would help them play. This underlying system of incomes and prices is something that the player could respond to in a “normal” way, that is, in the way economic theory predicts. Alternatively, because this is a fantasy environment in which no one really “dies”, the economic calculus may be unimportant and economic behavior correspondingly off-kilter.

3. Experimental Protocol

Research subjects were recruited from two undergraduate classes at Indiana University. Each participant received 20 points of course credit for playing the game and filling out a survey. Each student had to take a DVD with the game client on it, install it on their own computer, then create an account for online log-in.⁴ The system requirements were such that not all students in the two classes had a suitable computer at home. Students who could not participate or elected

⁴ The usual network architecture of a virtual world, which we followed, is that users have a piece of “client” software on their home computers and they use this client to communicate with the main game server, hosted by the game’s owners. The client is “dumb” in the sense that it knows nothing about the state of the virtual world; information about world state (location of characters, monster types and activities, location of treasure, quest states, etc.) is held secure at the central server. In our experiment, we hosted the world on our own server, and gave subjects client software that allowed them to log in to our server and experience the virtual world there.

not to participate were given an alternate assignment for equal credit. The final number of participants was 43.

When each participant was given login information, they received one of two IP addresses, one being “Arden A” and the other being “Arden B”. These identities (version A or B) were not revealed to the subjects; the subjects were not informed that there was more than one research environment. Subjects were given one or the other IP addresses based on the last digit of their Social Security number. Numbers 0-4 were assigned to Arden A (22 subjects), while 5-9 were assigned to Arden B (21 subjects).

Arden A was the working game environment developed during the 12-month production phase of the project. Arden B was created the day before the experiment began. The Arden A environment was loaded into our game editor, and then it was given a new name, Arden B, using the “Save As” feature of the editor. In other words, Arden B was an exact copy, bit for bit, of Arden A.

After Arden B was created, we then edited it as follows. In Arden A, there were four merchants, located in the early stages of the game, who sold Potions of Cure Light Wounds, for 15 gold pieces each. Each merchant had a stock of 1,000 potions, which we assumed (correctly) would far exceed the maximum possible demand. After creating Arden B, we edited these four merchants so that the price of a Cure Light Wounds Potion was 30 gold pieces. No other changes were made.

Once Arden A and B were created, they were installed on an internet server and assigned IP addresses. At this point, subjects were told that the server was “live” and that they could log in at will. We kept the servers “on” for one month (specifically, 32 days). As it happened, all of the participants had made their way through the content of the game by the 27th day. Thus, in closing the server, we did not truncate the game experience for any of the players. Moreover, since the experimental intervention involved Potions of Cure Light Wounds – an item that is useful only to characters in the early levels (levels 1 – 6 or so) – we can be certain that all of our subjects experienced the content that was relevant to our experiment (the early content), well before the experiment was closed.

Subjects were given brief written instructions on how to play the game. These instructions included, among many other things, the information that healing potions could be used to restore hit points during combat. The game client also allowed subjects to play through a tutorial, in the course of which they were introduced to merchants, gold pieces, and the mechanics of buying and selling.

Subjects obtained gold pieces primarily through the Dropped Loot items, which were placed in the streets and tunnels throughout the game world, in places where they were not difficult to spot. Specifically, there were 114 treasure items in the “Low” category (giving anything from 5 gold pieces to 16 gold pieces), 139

in the “Medium” category (2 – 32 gold), 18 in the “High” category (4 – 64 gold), and 11 in the “Wild” category (8 – 160 gold). A subject who found each of these treasures would, on average, acquire 5,096 gold pieces. Additional funds could be obtained by doing extra quests to obtain up to 5 magic items, which could then be sold to merchants for prices ranging from 100 gold to more than 2,000 gold. Alternatively, players could keep those items and use them. Just how much gold a player obtained was therefore dependent on the extent to which he pursued treasure, as well as the luck of the die. Unfortunately, due to the informational structure of our game-building software, we could not observe the gold accumulation of players. However, the amount of gold available seems sufficient to allow a typical player to purchase a number of Potions of Cure Light Wounds (whether priced at 15 gold or 30 gold).

In the course of the experiment, we discovered two bugs that had to be patched to allow two players to continue through the game. The bugs were of the nature of doors not opening even though the subject had the required key. Patching them required minor modification to the key-checking code. Both bugs occurred relatively late in the game content, and did not affect the health status or combat readiness of characters. Moreover, the bugs were exactly the same on both versions of Arden, and both were patched in exactly the same way. Thus it seems unlikely that the presence of these bugs had an impact on cross-world differences in purchases of Cure Light Wounds Potions.⁵

4. Results

Table 1 compares the two subject groups on numerous demographic lines. It would appear that the random assignment protocol was successful in creating two similar populations for the two environments. That the sample is overwhelmingly male is a product of the classes from which the sample was drawn: one was a class in The Videogame Industry, the other in Multiplayer Game Design. It is typical that most enrolled students in these two classes are male.

⁵ Patching in virtual worlds is unavoidable. It is to expected that some bugs will be discovered only after the world has gone live. With the world being always “on”, it will then be necessary to fix things “on the fly”. This led us to an important lesson for virtual world experiments: An experimental protocol cannot reasonably require that there be no changes “on the fly” during an ongoing experiment; there will always be bugs, and they will need to be patched. Rather, researchers should always report patches so that readers can review whether the patch was relevant to the research findings.

	Arden A (Control)	Arden B (Experiment)
Subjects	22	21
Age	22.7	22.1
Percent Male	87	93
Computer use hours per week	34.9	29.2
Percent with fathers who have a college degree or more	50.0	58.8

Table 1. Experimental and Control Demographics

The results of the experiment are summarized in Table 2. A reasonable number of potions were purchased in both worlds, 597 in the low-cost environment and 324 in the high-cost environment. Thus, it appears that the subjects found the potions to be useful in pursuing their goals within the game world.

	Arden A (Control)	Arden B (Experiment)
Subjects	22	21
Price of Cure Light Wounds Potion	15 gold	30 gold
Potions Purchased	597	324
Potions Per Subject	27.14	15.43

Table 2. Experimental Results

It also appears to be the case that the subjects were respondent to price incentives, as the Law of Demand suggests. Subjects in the low-cost world bought 27.14 potions on average, while those in the high-cost world bought only 15.43 Potions.

We judge this difference in potion consumption – a drop of almost half - to be substantively significant. However, because of a quirk in the game software that prevents us from calculating the sample variance of purchases, we are unable to conduct a full test of the statistical significance of the difference in means.⁶

⁶ To gather data, we placed a specific number of potions on potion merchants and then checked periodically how many were still there. The difference, we knew, must have been sold to players. However, this did not tell us who bought them. This is an example of a general problem with game-making software: the toolmakers give hardly any consideration to tools for data collection and analysis. It is very easy with contemporary tools to place a gnome in a cavern and have him interact in all kinds of ways with players. But it is almost impossible to get the gnome to record

Nonetheless, we are able to determine that under reasonable assumptions about the variance, the difference in sample means is statistically significant.⁷

We can use the basic results to estimate the Elasticity of Demand, a general statistic used for comparisons in economic studies of consumer behavior. According to the data, doubling the price of potions from 15 gold to 30 gold reduced quantity demanded by 43.1 percent:

$$100 \times [1 - (15.43 / 27.14)] = 43.1$$

The formula for Elasticity of Demand is

$$\epsilon = \text{Percent Change in Quantity Demanded} / \text{Percent Change in Price}$$

which in this case becomes

$$\epsilon = -43.1 / 100 = -0.431$$

A demand elasticity of -0.431 is well within the range of normal demand elasticities, that is, demand elasticities that one usually sees in studies of the real world.⁸ It thus appears that our research subjects' behavior did indeed conform to the Law of Demand.

what he has done and write the information to a database in a way that makes sense for quantitative analysis. Our pre-research reviews of world-building toolkits indicated that the Aurora Toolkit had the best chance of allowing the data collection we would need. We made an intense effort, over several months, to modify the Aurora Toolkit so that it would report auction-based market purchases and sales out to a database. Unfortunately, we were not able to build a reliable system of markets and reporting within the 15-month time frame of our grant. In the interests of time, we then moved to a protocol of basic observation, manually counting what had been sold.

⁷ As indicated in note 6, the game did not allow us to record how many potions each user purchased. Thus, while we can calculate the average potions per user as total potions divided by total users, we cannot calculate the standard deviation of potions purchased across users (since that would require the purchases of individuals). Some measurement of the standard deviation would of course be required for Student's t. If we assume that the standard deviation is equal to one-half of the mean (ie, that the lower tail of the distribution of purchases just includes 0), Student's t is 3.4981, which yields statistical significance at significance levels well below 5%. After exploring several assumptions about this variance, our sense is that only a very high variance would yield statistical insignificance; the difference in means here is simply very large relative to the means themselves.

⁸ To make this more concrete: We would be very surprised at a positive demand elasticity, since that would directly violate the Law of Demand. We would also be surprised at anything very large in absolute value (above 50). A good with demand elasticity of -10 would be considered very elastic. Larger elasticities approach a condition known as "perfectly elastic demand". Similarly, an elasticity of 0 would be termed "perfectly inelastic demand." Both of these conditions are considered empirical rarities. The vast majority of real-world demand elasticity estimates that the lead author has seen in his career fall between -2.0 and -0.05. The experimental finding is in this range.

5. Discussion

We have provided evidence that fantasy game players respond normally to basic economic incentives. Despite being immersed in a synthetic fantasy environment, in which there was no real reason to follow the usual rules of social and economic conduct, the research subjects in this study nonetheless purchased fewer health potions when those potions were more expensive. We estimate that the elasticity of demand for health potions in our game was -0.431 – a thoroughly normal demand elasticity.

Numerous caveats must be applied to this finding. First, this experiment was not as controlled as is usual in the existing field of social science experimentation. In a typical laboratory experiment in economics, political science, game theory, or other fields, a similar number of subjects (20 – 50) spend several hours at once in a computer lab. Their communications are closely monitored and controlled. The information needed for their participation is presented in a strict protocol involving the reading of instructions and/or placement on a computer screen. Experimenters tightly restrict the narrative and contexts, to prevent subjects from inventing a “story” within which to place their decision-making. Whatever relationships, social or political or economic, there may be among the participants are either created by the experimenters or, if not, tightly controlled and sanitized by them for the duration of the test.

Our experiment, by contrast, did not occur in a lab. Communication was not controlled or observed. Instructions were provided in writing but no effort was made to ensure that subjects read the instructions or interpreted them in the same way. Most strikingly, our experiment embedded decision-making in a wild and detailed narrative, a rich context that was not only not sterile but fantastically different from the context of the real world. This was, of course, by design, in order to test whether such a thing made a difference in our theories of human behavior. We found in this case that it did not. We cannot conclude from this, however, that the narrative context will never make a difference. Rather, we are simply presenting evidence that one can give up the high levels of control that are possible in the laboratory, in order to achieve other advantages in terms of experimental scope in time and population. Nonetheless, the relative lack of experimental control, when compared to laboratory experiments, is a dimension along which studies like this can be legitimately criticized. In future work, protocols should be explored by which the communications of participants can be observed or regulated; and future experiments must determine where context matters and where it does not.

A second caveat might be that this finding is not as novel as we claim it to be. Numerous researchers have discovered that “normal” economic behavior occurs in many places where one might not expect it, such as POW camps (Radford, 1945), mental wards (Battalio et al., 1974; Tarr, 1976; Battalio et al., 1977) and among pigeons (Battalio, 1981). Moreover, experimental economists have for a

long time used linked computer networks to do experiments and to teach students (Williams and Walker, 1993; see again Chesney, Chuah, and Hoffman, 2007). Thus, our study only contributes the comparatively minor piece of information that one aspect of economically normal behavior – price sensitivity – occurs in online fantasy games. Some readers, familiar with these other studies, will find that to be unsurprising. However, we argue that the main question here is not about rationality per se, or about online networks per se, but rather about the effect of fantasy. It is not clear a priori whether, once gamers enter the fantasy zone, they become more like rational pigeons, or more like irrational pixies. Prior to embarking on this search, we encountered many skeptics, in many different fields, who argued that the fantasy context must invalidate most human behavioral theory. We have found evidence to the contrary. We feel this finding is relevant for numerous discussions, in numerous fields.

A third caveat worth noting is that our study was based on a few students, mostly male, at a large Midwestern university. It was not possible to observe the variance of purchases, and therefore tests of statistical significance were incomplete. Future work should take advantage of the flexibility of this technology and conduct virtual world experiments with many more people, from many more places, with better measurements. Some of this, of course, requires the relaxation of control, the subject of our first caveat. It would thus seem that there is a tradeoff between control and generality. Perhaps future methodological work could explicate this tradeoff and seek to reduce it. In any case, one great advantage of our protocol is that it is trivial to replicate. Future researchers could simply run the experiment with more subjects and better observational tools, to see whether our finding is unusual or not.⁹

Despite these clear limitations of the study, we argue that it is a reasonably solid first step in a promising new direction. There are major methodological advantages to addressing macro-scale social science questions using virtual world petri dishes. These advantages are shared with any experimental approach. We know, for example, that we have identified a change in quantity demanded – a movement along the demand curve - because of the experimental condition. The demand curve indicates the quantity demanded at a given price. It shows how the amount purchased should change as the price of the good changes, all else equal. In the real world, and with any kind of historical data, identifying movements along the demand curve is difficult. We observe a price change, and a change in the quantity sold, but we can never be sure that “all else” is truly “equal.” On the contrary, with real-world data, we can be fairly confident that nothing else has remained fixed. Many other things have changed as well. The price of oranges rose, but something also probably happened to the price of grapefruits, the price of gasoline, and the incomes of people who drink juice for breakfast. Because of these other factors, we cannot be sure that the change in orange sales following a change in orange prices is truly a trace along

⁹ The experimental environment is available for free download at swi.indiana.edu. Follow the instructions there to run precisely the same experiment that we did.

the demand curve for oranges. Here, by contrast, we can be reasonably sure that we have identified a trace along the demand curve, because we can be reasonably sure that all else really is equal. The two worlds we have examined differed only in price. The populations of those worlds were apparently composed of very similar people, and nothing was done to one group that was not done to the other. Thus we can be reasonably confident that the demand elasticity of Potions in our world is actually -0.431 at the current market point.

This degree of confidence in identifying causality is one of the great advantages of the experimental method in general. Of particular interest here, however, is the possibility of conducting experiments on a much larger scale, both in terms of time and population. Game environments seem particularly attractive as forms of entertainment, and it has been observed that some game environments have millions of players for many years. Can these huge worlds be used to study human behavior on a vast scale? This study suggests indeed that the presence of fantastical creatures and odd roles within these environments need not invalidate them as research sites. It may be true that players enter a game world only because they want to hunt dragons; but if they need health potions to do this, we can expect that the market for those potions will be, at base, a pretty normal market, with ordinary supply and demand curves and ordinary equilibrium conditions. This at least is the conclusion toward which the research in this paper points. Determining how far this kind of generality extends, however, is a task that should occupy the research agenda in this area for many years.

References

- Balicer, R., "Modeling Infectious Diseases Dissemination Through Online Role-Playing Games," *Epidemiology* 18(2), pp 260-261.
- Battalio, Raymond C., Edwin B. Fisher, Jr., John H. Kagel, Robert L. Basmann, Robin C. Winkler, Leonard Krasner (1974), "An Experimental Investigation of Consumer Behavior in a Controlled Environment," *Journal of Consumer Research* 1(2), pp. 52-60.
- _____, John H. Kagel, and Morgan O. Reynolds (1977), "Income Distributions in Two Experimental Economies," *Journal of Political Economy* 85(6), pp. 1259-1271.
- _____, John H. Kagel, Howard Rachlin and Leonard Green (1981), "Commodity-Choice Behavior With Pigeons as Subjects," *Journal of Political Economy*, 89(1), pp. 67-91.
- Castronova, Edward (2001), "Virtual Worlds: A First-Hand Account of Market and Society on the Cyberian Frontier," CESifo Working Paper No. 618, December 2001.
- _____, (2005), *Synthetic Worlds: The Business and Culture of Online Games*, Chicago: University of Chicago Press.
- _____, (2006), "On the Research Value of Large Games: Natural Experiments in Norrath and Camelot." *Games and Culture*. Vol. 1(2), pp 163-186.
- _____, (2007), *Exodus to the Virtual World: How Online Fun is Changing Reality*, New York: Palgrave/MacMillan.
- Chesney, Thomas, Swee-Hoon Chuah, and Robert Hoffman (2007) "Virtual World Experimentation: An Exploratory Study," *Nottingham University Business School Occasional Paper Series*, No. 2007-21.
- Jakobsson, Mikael and T.L. Taylor (2003), "The Sopranos Meets EverQuest: Socialization Processes in Massively Multi-User games," *FineArt Forum* 17(8),
http://www.fineartforum.org/Backissues/Vol_17/faf_v17_n08/reviews/jakobsson.html.
- Lastowka, F. Greg and Dan Hunter (2004), "The Laws of the Virtual Worlds," *California Law Review* 92(1).

- Lofgren, E.T. and N. H. Fefferman (2007), "The Untapped Potential of Virtual Game Worlds to Shed Light on Real World Epidemics," *Lancet Infectious Diseases* 2007(7), pp 625-629.
- Lucas, Jr. Robert E. (1986), "The Behavioral Foundations of Economic Theory," *The Journal of Business*, 59(4), pp. S401-S426.
- Malaby, Thomas H. (2006), "Parlaying Value: Capital in and Beyond Virtual Worlds," *Games and Culture* 1(2), pp. 141-162.
- Radford, R. A. (1945), "The Economic Organization of a P.O.W. Camp," *Econometrica* 12(48), pp. 189-201.
- Reeves, Byron and Clifford Nass (1996), *The Media Equation: How People Treat Computers, Television, and New Media like Real People and Places*, Cambridge: Cambridge University Press.
- Tarr, David G. (1976), "Experiments in Token Economies: A Review of the Evidence Relating to Assumptions and Implications of Economic Theory," *Southern Economic Journal*, 43(2) pp. 1136-1143.
- Williams, Arlington W. and James M. Walker (1993), "Computerized Laboratory Exercises for Microeconomics Education: Three Applications Motivated by Experimental Economics," *Journal of Economic Education* 24(4), pp. 291-315.
- Yee, Nick and Jeremy Bailenson (2007), "The Proteus Effect: The Effect of Transformed Self-Representation on Behavior," *Human Communication Research* 33, pp. 271-290.
- Zimbardo, P. G. (2007), *The Lucifer Effect: Understanding How Good People Turn Evil*. New York: Random House.

CESifo Working Paper Series

for full list see www.cesifo-group.org/wp

(address: Poschingerstr. 5, 81679 Munich, Germany, office@cesifo.de)

- 2289 Mikael Priks, Do Surveillance Cameras Affect Unruly Behavior? A Close Look at Grandstands, April 2008
- 2290 Marianna Belloc and Daniela Federici, A Two-Country NATREX Model for the Euro/Dollar, April 2008
- 2291 Nicolas Treich, The Value of a Statistical Life under Ambiguity Aversion, April 2008
- 2292 J. Atsu Amegashie, Socially-Tolerable Discrimination, April 2008
- 2293 M. Hashem Pesaran and Andreas Pick, Forecasting Random Walks Under Drift Instability, April 2008
- 2294 Steven Brakman, Gus Garita, Harry Garretsen and Charles van Marrewijk, Unlocking the Value of Cross-Border Mergers and Acquisitions, May 2008
- 2295 Eric O'N. Fisher and Kathryn G. Marshall, The Structure of the American Economy, May 2008
- 2296 Claudia M. Buch and Martin Schlotter, Regional Origins of Employment Volatility: Evidence from German States, May 2008
- 2297 Helmuth Cremer, Philippe De Donder, Dario Maldonado and Pierre Pestieau, Taxing Sin Goods and Subsidizing Health Care, May 2008
- 2298 Reinhilde Veugelers and Frederick van der Ploeg, Reforming European Universities: Scope for an Evidence-Based Process, May 2008
- 2299 Jon H. Fiva and Lars J. Kirkebøen, Does the Housing Market React to New Information on School Quality?, May 2008
- 2300 Tina Klautke and Alfons J. Weichenrieder, Interest Income Tax Evasion, the EU Savings Directive, and Capital Market Effects, May 2008
- 2301 Harald Badinger and Peter Egger, GM Estimation of Higher Order Spatial Autoregressive Processes in Panel Data Error Component Models, May 2008
- 2302 Jan K. Brueckner, Slot-Based Approaches to Airport Congestion Management, May 2008
- 2303 Sören Blomquist, Vidar Christiansen and Luca Micheletto, Public Provision of Private Goods and Nondistortionary Marginal Tax Rates, May 2008
- 2304 Dan Anderberg and Alessandro Balestrino, The Political Economy of Post-Compulsory Education Policy with Endogenous Credit Constraints, May 2008

- 2305 Tomer Blumkin, Yoram Margalioth and Efraim Sadka, The Role of Stigma in the Design of Welfare Programs, May 2008
- 2306 Vesa Kannianen and Paolo M. Panteghini, Tax Neutrality: Illusion or Reality? The Case of Entrepreneurship, May 2008
- 2307 Thomas Dohmen, Armin Falk, David Huffman and Uwe Sunde, The Intergenerational Transmission of Risk and Trust Attitudes, May 2008
- 2308 Guglielmo Maria Caporale and Mario Cerrato, Using Chebyshev Polynomials to Approximate Partial Differential Equations, May 2008
- 2309 Peter Egger and Doina Maria Radulescu, Labour Taxation and Foreign Direct Investment, May 2008
- 2310 Laurent Linnemer, Dissipative Advertising Signals Quality even without Repeat Purchases, May 2008
- 2311 Jordi Jofre-Monseny and Albert Solé-Ollé, Which Communities should be afraid of Mobility? The Effects of Agglomeration Economies on the Sensitivity of Firm Location to Local Taxes, May 2008
- 2312 Andreas Haufler and Ferdinand Mittermaier, Unionisation Triggers Tax Incentives to Attract Foreign Direct Investment, May 2008
- 2313 Ronel Elul and Piero Gottardi, Bankruptcy: Is it enough to Forgive or must we also Forget?, May 2008
- 2314 Andreas Irmen and Johanna Kuehnel, Productive Government Expenditure and Economic Growth, May 2008
- 2315 Beate Henschel, Carsten Pohl and Marcel Thum, Demographic Change and Regional Labour Markets: The Case of Eastern Germany, May 2008
- 2316 Gabriel Felbermayr, Wido Geis and Wilhelm Kohler, Restrictive Immigration Policy in Germany: Pains and Gains Foregone?, May 2008
- 2317 Michael Hofmann, Gerhard Kempkes and Helmut Seitz, Demographic Change and Public Sector Budgets in a Federal System, May 2008
- 2318 Paul De Grauwe, Macroeconomic Modeling when Agents are Imperfectly Informed, June 2008
- 2319 Johann K. Brunner and Susanne Pech, Optimum Taxation of Inheritances, June 2008
- 2320 Thomas Eichner and Marco Runkel, Corporate Income Taxation of Multinationals in a General Equilibrium Model, June 2008
- 2321 Rainald Borck and Matthias Wrede, Subsidies for Intracity and Intercity Commuting, June 2008

- 2322 Patricia Apps and Ray Rees, Testing the Pareto Efficiency of Household Resource Allocations, June 2008
- 2323 Amihai Glazer, Vesa Kannianen and Panu Poutvaara, Firms' Ethics, Consumer Boycotts, and Signalling, June 2008
- 2324 Claudia M. Buch, Jörg Döpke and Kerstin Stahn, Great Moderation at the Firm Level? Unconditional vs. Conditional Output Volatility, June 2008
- 2325 Helmuth Cremer, Philippe De Donder, Dario Maldonado and Pierre Pestieau, Forced Saving, Redistribution and Nonlinear Social Security Schemes, June 2008
- 2326 M. Hashem Pesaran and Paolo Zaffaroni, Optimal Asset Allocation with Factor Models for Large Portfolios, June 2008
- 2327 Harald Badinger and Peter Egger, Horizontal versus Vertical Interdependence in Multinational Activity, June 2008
- 2328 Jan K. Brueckner and Harris Selod, A Theory of Urban Squatting and Land-Tenure Formalization in Developing Countries, June 2008
- 2329 Paolo M. Panteghini, Corporate Debt, Hybrid Securities and the Effective Tax Rate, June 2008
- 2330 Guglielmo Maria Caporale, Juncal Cuñado and Luis A. Gil-Alana, Modelling Long-Run Trends and Cycles in Financial Time Series Data, June 2008
- 2331 Avi Ben-Bassat and Momi Dahan, Social Identity and Voter Turnout, June 2008
- 2332 Martin R. West and Ludger Wößmann, "Every Catholic Child in a Catholic School": Historical Resistance to State Schooling, Contemporary Private Competition, and Student Achievement across Countries, June 2008
- 2333 Erkki Koskela and Panu Poutvaara, Outsourcing and Labor Taxation in Dual Labor Markets, June 2008
- 2334 Philippe Choné and Laurent Linnemer, Optimal Litigation Strategies with Signaling and Screening, June 2008
- 2335 Albert Solé-Ollé and Pilar Sorribas-Navarro, Does Partisan Alignment Affect the Electoral Reward of Intergovernmental Transfers?, June 2008
- 2336 Antonio Cabrales and Piero Gottardi, Markets for Information: Of Inefficient Firewalls and Efficient Monopolies, June 2008
- 2337 Sumon Majumdar and Sharun W. Mukand, The Leader as Catalyst – on Leadership and the Mechanics of Institutional Change, June 2008
- 2338 Ulrich Hange, Tax Competition, Elastic Labor Supply, and Growth, June 2008

- 2339 Guy Laroque and Bernard Salanié, Does Fertility Respond to Financial Incentives?, June 2008
- 2340 Adriano Paggiaro, Enrico Rettore and Ugo Trivellato, The Effect of Extending the Duration of Eligibility in an Italian Labour Market Programme for Dismissed Workers, June 2008
- 2341 Helmut Seitz, Minimum Standards, Fixed Costs and Taxing Autonomy of Subnational Governments, June 2008
- 2342 Robert S. Chirinko, Leo de Haan and Elmer Sterken, Asset Price Shocks, Real Expenditures, and Financial Structure: A Multi-Country Analysis, July 2008
- 2343 Wolfgang Leininger, Evolutionarily Stable Preferences in Contests, July 2008
- 2344 Hartmut Egger and Udo Kreickemeier, Fairness, Trade, and Inequality, July 2008
- 2345 Ngo Van Long and Bodhisattva Sengupta, Yardstick Competition, Corruption, and Electoral Incentives, July 2008
- 2346 Florian Baumann, Employment Protection: The Case of Limited Enforceability, July 2008
- 2347 Alessandro Balestrino, Cinzia Ciardi and Claudio Mammini, On the Causes and Consequences of Divorce, July 2008
- 2348 Dirk Schindler and Benjamin Weigert, Insuring Educational Risk: Opportunities versus Income, July 2008
- 2349 Lammertjan Dam and Ben J. Heijdra, The Environmental and Macroeconomic Effects of Socially Responsible Investment, July 2008
- 2350 Avner Greif, Contract Enforcement and Institutions among the Maghribi Traders: Refuting Edwards and Ogilvie, July 2008
- 2351 Helmuth Cremer, Philippe De Donder, Dario Maldonado and Pierre Pestieau, Habit Formation and Labor Supply, July 2008
- 2352 Francesco Menoncin and Paolo M. Panteghini, The Johansson-Samuelson Theorem in General Equilibrium: A Rebuttal, July 2008
- 2353 Michael Kaganovich and Itzhak Zilcha, Alternative Social Security Systems and Growth, July 2008
- 2354 Keith Blackburn, Kyriakos C. Neanidis and M. Emranul Haque, Corruption, Seigniorage and Growth: Theory and Evidence, July 2008
- 2355 Edward Castronova, A Test of the Law of Demand in a Virtual World: Exploring the Petri Dish Approach to Social Science, July 2008