

Castronova, Edward; Falk, Matthew

Working Paper

Virtual Worlds as Petri Dishes for the Social and Behavioral Sciences

RatSWD Working Paper, No. 47

Provided in Cooperation with:

German Data Forum (RatSWD)

Suggested Citation: Castronova, Edward; Falk, Matthew (2008) : Virtual Worlds as Petri Dishes for the Social and Behavioral Sciences, RatSWD Working Paper, No. 47, Rat für Sozial- und Wirtschaftsdaten (RatSWD), Berlin

This Version is available at:

<https://hdl.handle.net/10419/186163>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



German Council for Social
and Economic Data (RatSWD)

www.ratswd.de

RatSWD

Working Paper Series

Working Paper

No. 47

Virtual Worlds as Petri Dishes for the Social and Behavioral Sciences

Edward Castronova, Matthew Falk

December 2008



Federal Ministry
of Education
and Research

Working Paper Series of the Council for Social and Economic Data (RatSWD)

The *RatSWD Working Papers* series was launched at the end of 2007. Since 2009, the series has been publishing exclusively conceptual and historical works dealing with the organization of the German statistical infrastructure and research infrastructure in the social, behavioral, and economic sciences. Papers that have appeared in the series deal primarily with the organization of Germany's official statistical system, government agency research, and academic research infrastructure, as well as directly with the work of the RatSWD. Papers addressing the aforementioned topics in other countries as well as supranational aspects are particularly welcome.

RatSWD Working Papers are non-exclusive, which means that there is nothing to prevent you from publishing your work in another venue as well: all papers can and should also appear in professionally, institutionally, and locally specialized journals. The *RatSWD Working Papers* are not available in bookstores but can be ordered online through the RatSWD.

In order to make the series more accessible to readers not fluent in German, the English section of the *RatSWD Working Papers* website presents only those papers published in English, while the German section lists the complete contents of all issues in the series in chronological order.

Starting in 2009, some of the empirical research papers that originally appeared in the *RatSWD Working Papers* series will be published in the series *RatSWD Research Notes*.

The views expressed in the *RatSWD Working Papers* are exclusively the opinions of their authors and not those of the RatSWD.

The RatSWD Working Paper Series is edited by:

Chair of the RatSWD (2007/ 2008 Heike Solga; 2009 Gert G. Wagner)

Managing Director of the RatSWD (Denis Huschka)

Virtual Worlds as Petri Dishes for the Social and Behavioral Sciences

By Edward Castronova and Matthew Falk

Indiana University, Department of Telecommunications

December 2008

Abstract

The next tool for social science experimentation should allow for macro level, generalizable, scientific research. In the past devices such as rat mazes, Petri dishes and supercolliders have been developed when scientists needed new tools to do research. We believe that Virtual Worlds are the modern equivalent to supercolliders for social scientists, and feel they should be the next area to receive significant attention and funding. The advantages provided by virtual worlds research outweigh the costs. Virtual worlds allow for societal level research with no harm to humans, large numbers of experiments and participants, and make long term and panel studies possible. Virtual worlds do have some drawbacks, in that they are expensive and time consuming to build. These obstacles can be overcome, however, by adopting the models of revenue and maintenance practiced by the current game industry. The returns from virtual worlds being used as scientific tools could reach levels that would self fund future research for decades to come. However, at the beginning an investment of funding agencies seems to be necessary.

Keywords: Virtual Worlds, Macro Level Experiments, Research Infrastructure

JEL Classification: C15, C59, C82, C99

1 Introduction

In the past, science has developed new tools for research as the need arose. From “Petri Dishes” to “Rat Mazes”, and on to even “Supercolliders”, scientists require specific tools to answer the questions they ask. These devices are all influential for specific micro-level observations. But when it comes to social science and societal level research questions the tools for empirical research were not developed much further for almost two centuries. Further developments occurred in the collection of survey data after World War II. And in the nineties “experimental economics” started to become a new and popular tool for empirical research (which is surprisingly seldom applied by sociologists). Now social scientists should be looking to a new place: Virtual Worlds (VW).¹

To be considered a virtual world, a game or social networking site must: be computer generated, persistent (always there even if no one is currently logged into it), and have humans represented in the form of Avatars (the embodiment of the user in the virtual space) capable of taking actions on behalf of their human counterpart. Only in virtual worlds do we find the proper tool set for large scale social science research, previously unavailable to scientists. These defining features combine to allow scientists access to long running persistent societies of users, all engaged in actions that resemble what we see in the real world (Castronova and Falk, Forthcoming)². Please note that we talk in this paper about experiments on the macro level of a society. We do not address the issue that one can use virtual worlds for running experiments on the micro level of individual players. For this possibility see e. g. Chesney, Chuah and Hoffmann (2007). For an even simpler approach which uses virtual worlds, like “Second Life”, for social surveys see Bell, Castronova and Wagner (2008).

Because of the large scale commercial success and now wide spread use of VW’s it is possible to collect large amounts of data from large numbers of users. Instead of a few hundred people in one place for a short time, as in current, e. g., experimental economics and other lab based research, we can draw from populations that range from thousands to millions and take measurements over time. Because of the size of the populations, VW’s let us look at causation at the macro or societal level.

Virtual worlds range in scope from small scale, internet browser based games with perhaps a few hundred players to the massively successful game World of Warcraft, which has had ten million subscriptions purchased in its four years of existence, and has an estimated consistent player base of 8 million. The populations of these worlds span the globe, and it is just as possible to meet someone from thousands of miles away as it is to join your friends from down the street when exploring the virtual world. (Castronova 2005, 2007)

2 Petri Dishes, Rat Mazes, Supercolliders

While the virtual world is not a sealed vacuum, it does resemble a Petri dish in its functionality (Castronova and Falk, Forthcoming). Many users, millions at a time in fact, can exist in a game like World of Warcraft. These users are not, however, all interacting with each other in one space. Like Petri dishes in the laboratory, individual servers – digital copies of the same world with unique users interacting – make it possible for technology to handle the demand. The servers, or individual Petri dishes, contain the same ingredients in them. It is the bacteria, or in this case the players, that differ based on which server they choose to play.

Because the servers all start inherently as exactly the same world, it is possible to make one small change to the composition of the goo in the Petri dish – a single variable on one server – and create experimental conditions. Server after server, side by side, can resemble rows of Petri dishes in a lab. One group containing a set of control conditions, another group one experimental condition, and so on and so forth; the underlying code, or the thing that makes it all work, does not change. The color of the sky, the names of the places, and the size of any oceans do not differ, unless of course that is what the scientist chooses to change. As will be addressed later, the only restriction to the number of servers and amount of players is monetary.

1 This paper is written within and for the project „Developing the Research Data Infrastructure for the Social and Behavioral Sciences in Germany and Beyond: Progress since 2001, Current Situation, and Future Demands“ of the German Council for Social and Economic Data (RatsWD) (<http://www.ratswd.de/ver/veranstaltungen.php>).


2 See also Giles (2007) and Miller (2007).

Just because this is possible does not make virtual worlds the perfect answer for all questions. In fact, there are some types of questions that virtual worlds are poor at answering. Much like all experimental tools, the tool must be designed to answer the types of questions that the researcher wants to ask.

For instance, mammalian cognition is a very often studied topic. Some scientists use rat mazes to test the cognitive habits of rats, and others recreate the mediated environments humans encounter every day to examine what effects this has on the brain. Both of these cases provide good examples of how scientific tools are shaped to the questions at hand.

Figure 1 demonstrates the relationship between type of question asked and the tools used to study them:

Figure 1. Experimental Tools and Scientific Questions



Source: Castronova and Falk (forthcoming).

The horizontal axis in Figure 1 arrays experimental environments according to their fidelity to reality. An environment that is very concrete replicates reality quite well. It is a simulation. A media effects environment that puts a TV with currently running programming, in an American living room, replicated right down to the six-pack and the cat odor, is concrete. A media effects environment that attaches wires to individual's heads and has them watch triangles on a small screen while holding a buzzer, is abstract.

The diagonal dashed line in the figure illustrates the set of ideal experiments, ones where the conceptual level of the question is well-matched by the concreteness of the experimental environment. If the research question is specific, the experimental environment must be concrete. If the question is general, the experimental environment must be abstract. It would be a bad experiment (labeled "Bad Experiments" in the figure) to try to study a specific question with an abstract research environment. You cannot conclude much about the reaction of typical American families to last night's newscast by wiring their heads and asking them to watch triangles on a screen. Bad experiments can go the other way, too: you cannot learn much about the general rate of response times to visual stimulus just by watching people in their living rooms.

The point of the diagram is this: it is senseless to make claims about the validity of a research

environment unless you know what sort of question is being studied. A rat maze is a terribly abstract environment; yet would anyone say “You can’t learn anything about anything in a rat maze. A rat maze is too unlike the real world.” It is possible to learn a great deal about mammalian cognition this way. When the research questions involve societies, or have macro level implications, we must build a more concrete and specific environment.

3 Current State of the Research Field

While many researchers examine communications and media in the form of Virtual Worlds, only a very small few are using empirical methods. Much of the research is concerned with theory creation, observational and ethnographic methods, including observational analyses by means of regression analysis. All of these methods are valid, and collect pieces of information, but none of them are experimental and as such do not lead to the concrete, generalizable, macro level information about human behavior that social science is seeking.

Social and behavioral scientists seek to understand how humans interact. Social scientists, specifically, want to explore the large questions of human interactions: war, disease, starvation, ecological disaster, economic stability, etc. The only way to solve macro level problems is with macro level science. *Controlled experimentation at the societal level* is not being conducted, and in fact would be impossible to conduct under normal circumstances. It is inadvisable and indeed impossible to remove real humans from society, place them in a vacuum for months or years at a time and then experiment on them. There are some attempts being made at this type of science, known to economists as “field experiments” (Harrison and List, 2004. Hausman and Wise, 1985. List, 2008) but they are often derided for their inability to produce controlled and generalized results.

Field experiments take on one of three forms, artificial field experiments, framed field experiments, and natural field experiments. Artificial field experiments tend to resemble laboratory experiments as closely as possible, but with a sample drawn from a specific population of interest thus eliminating generalizability, unless that group is itself representative of the general population. Framed field experiments entail placing experimental differences in their natural habitat, such as differing social programs for groups and devising which choice worked better overall. These again target a specific population (those participating in that particular choice at that location at that time) and while less abstract than artificial field experiments, do not hold up to the rigor of laboratory testing standards. The closest of the three to laboratory science is natural field experiments, which combines the anonymity on the part of the subject with the experimental manipulation of framed field experimentation. They fall short again, however, due to the interaction of natural environment, and the lack of availability of opportunity to produce works; researchers are limited to relying on the presence naturally occurring phenomena that they can gain approval to study.

Governmental interest in research falls mostly into a slightly different category of its own: simulation research. On the surface, simulation research looks much like the virtual worlds research we propose here, but there is a fundamental difference between the two. Simulations are essentially computer run models, in which the players (known as “agents”) are also computerized. Each behaves in a manner that is simple and is predictive of how an individual would act, given that each individual would continue to always make the most “rational” choice. This is problematic, as many believe that humans do not react rationally to many if any situations, and therefore consider the interactions of simple, rational models to be incomplete. A fundamental improvement to the model would be to use real humans in place of the agents – exactly what virtual worlds offer.

Building these social science equivalents to a supercollider appears both necessary, and expensive. Preliminary forays into this research field are already being conducted. Our group, the Synthetic Worlds Initiative of Indiana University, has already completed the building of a small scale virtual world, and the subsequent experimentation process within.

Based on a \$250,000 dollar grant from the MacArthur Foundation, “Arden” – a world based loosely on the works of William Shakespeare – took a student team almost two years to build. It then took another several months to run an experiment within the world, and compile those results for publication (Castronova, 2008). The experimental test run in Arden was a check of the economic law of demand. Having found that the law of demand holds true in a virtual world, Arden was deemed a successful first step towards the virtual Petri dish, or supercollider. But it was, in fact, only a first step.

The next logical step is to create another virtual world, capable of housing more users, and answering bigger questions. Along these lines, we are currently developing a game called Greenland,

which will be used to test the emergence of currency in the form of a web browser based resource collection game. But even this project, which can expect somewhere between several hundred and several thousand subjects is merely another small development. The ideal supercollider-level virtual world would be more like World of Warcraft, and consequently cost much more.

4 Development Costs, Future Research and Recommendations

Developing a persistent virtual world game of A-list, or top quality, requires not only a significant investment of time and personnel, but also has a large overhead cost for start up. This can be an insurmountable obstacle in terms of current social science research funding awards. Other areas of research and public service provide examples of how to research and develop massively expensive projects that get results and, in the end, generate profits to replenish the research and development costs.

Virtual world development costs are typically held in secrecy, as game design companies do not want to publicize exactly how much they've spent developing their projects. However, through the shared knowledge of former and current leaders in that business, it is quite possible to make inferences as to how much one should cost. Game development costs come primarily from three areas: the game design team / development, game launch, and customer service during the years in which a world operates.

Game design teams are typically small to start, possibly 5 or so people, but rapidly expand to include teams of 25-40 people, depending on the size of the project. This expansion occurs over a couple of years, and projects regularly take 30+ months from initial design meeting to the end of testing, or launch. For example, in 2005, a rule of thumb was that it cost approximately 10,000 USD per month per person on the team. This does not match up with the current size of research funding in the field – remember that Arden was created on a 250,000 dollar grant over the course of two years by a team of approximately 15 people. Professional game developers also work 40+ hour weeks, and are dedicated staff, whereas Arden and Greenland are being developed by graduate students working part time. Hiring professional staff would greatly speed up projects like this, and allow for faster game development and more experimentation.

Beyond the cost of personnel, the servers (*Petri dishes*) discussed earlier are each quite an expensive investment, and require a large amount of expensive bandwidth to run. Setting all this up and making sure it works before releasing it to the public is the next step in game development. In 2003, for example, the average for game “launch” (as it is known in the industry) was 7 million, with amounts in the 10-12 million range being more the rule than the exception. Current research is being performed on two small servers being hosted on university campuses. These servers simply cannot handle the mass amounts of players that the servers game companies use, and therefore limit the amount of study participants and instances of the world it is possible to have.

After the game goes public, and research and play commence, there is still a significant time and money investment required on the part of the game support staff. They must maintain player relations, collect subscription fees (if using one of the fee based models we will discuss below), and take care to maintain the software and hardware that allow users to access the world.

There are, however, two examples of ways in which it is possible to fund large public projects that subsequently pay the funding agencies back, and create profits to be used for further research/projects. These parallels can be drawn between nuclear power plants, the pharmaceutical research industry, and our vision of a virtual world as a supercollider.

Nuclear power plants, much like new experimental drugs, have been funded on public money to begin with. These infrastructures, once built, begin to sell their services to customers (in the form of power and pills, respectively). Through this revenue stream the companies that undertook the burden of building and maintaining the facilities (in the case of nuclear power) and developing, researching and testing (in the case of drugs) pay back the start up money it required to make those advances.

This is also the model used by the game industry around the world today. Games are launched to the public with both a “box fee” (the price the consumer pays in the store for the software) and then a monthly subscription fee. For example, upon release World of Warcraft cost approximately 50 USD, and since also costs users around 15 USD a month to play. These fees mean that Blizzard Entertainment, the parent company that financed the creation of the project, has seen its money back and more. Blizzard continues to use the profits from World of Warcraft to not only pay the aforementioned support staff, but to fund new projects as well. It is important to remember though,

that it does take time to see this return on investment. Typically 12 months, at the minimum, if a game is a large commercial success. If it is not, this process can take much longer. This does present a valid and established format for funding agencies to consider when making choices about funding large projects of this nature.

5 Conclusion

The next tool for social science experimentation should allow for macro level, generalizable, scientific research. In the past devices such as rat mazes, Petri dishes and supercolliders have been developed when scientists needed new tools to do research. We believe that Virtual Worlds are the modern equivalent to supercolliders for social scientists, and feel they should be the next area to receive significant attention and funding. The advantages provided by virtual worlds research outweigh the costs. Virtual worlds allow for societal level research with no harm to humans, large numbers of experiments and participants, and make long term and panel studies possible.

Virtual worlds do have some drawbacks, in that they are expensive and time consuming to build. These obstacles can be overcome, however, by adopting the models of revenue and maintenance practiced by the current game industry. The returns from virtual worlds being used as scientific tools could reach levels that would self fund future research for decades to come. However, at the beginning an investment of funding agencies seems to be necessary.

References

- Bell, M., E. Castronova, and G G. Wagner (2008). Virtual Assisted Self Interviewing (VASI): An Expansion of Survey Data Collection Methods to the Virtual Worlds by Means of VDCI. RatSWD Working Paper No. 42. Berlin. Available at http://www.ratswd.de/download/workingpapers2008/42_08.pdf
- Castronova, E. et al. (1998). A Test of the Law of Demand in a Virtual World: Exploring the Petri Dish Approach to Social Science. CESifo Working Paper Series No. 2355. Munich. Available at SSRN: <http://ssrn.com/abstract=1173642>
- Castronova, E. (2007). Exodus to the Virtual World How Online Fun is Changing Reality. New York, NY, Palgrave Macmillan.
- Castronova, E. (2005). Synthetic Worlds. Chicago, IL, The University of Chicago Press.
- Castronova, E and M. Falk. (Forthcoming). "Virtual Worlds: Petri Dishes, Rat Mazes, Supercolliders." Games and Culture Special Issue, Forthcoming.
- Chesney, Th., S.-H. Chuah, and R. Hoffmann. (2007). Virtual world experimentation: An exploratory study. *Nottingham University Business School - Industrial Economics Division Occasional Paper Series No. 2007-21*, Nottingham: <http://d.repec.org/n?u=RePEc:nub:occpap:21&r=exp>.
- Giles, J. (2007). Life's a Game. *Nature* 445, 18-20.
- Harrison, G.W. & List, J.A. (2004). Field experiments. *Journal of Economic Literature*, 42(4). Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=698961
- Hausman, J.A. & Wise, D.A. (1985). *Social Experimentation*. University of Chicago Press: Chicago.
- List, J. A. (2008). Homo experimentalis evolves. *Science*, 321, 207-208. Available at <http://www.sciencemag.org/cgi/reprint/321/5886/207.pdf>
- Miller, G. (2007). The Promise of Parallel Universes. *Science* 317, 1341-1343.