"For !!SCIENCE!!": Examining Epistemic Practices of the Community of Players of Dwarf Fortress

Mario Martinez-Garza, Vanderbilt University, Nashville, TN, USA

ABSTRACT

Dwarf Fortress is a digital game with an unusually detailed and complex underlying simulation based on real world systems. Gameplay is unforgiving and condensed knowledge resources that support effective play are scarce. To learn more about the inner workings of the game and make its challenges more tractable, the community of Dwarf Fortress players engages in systematic, evidence-based, experimental inquiry of the game. The community calls this pursuit "dwarf science". In this paper, the author investigates the origins, evolution, and practice of "dwarf science", and frames it as a model of how digital games for science learning might support the epistemic frame of science among learners.

Keywords: *Construction and Management Simulation (CMS), Dwarf Fortress, Dwarf Science, Inquiry-based Science Classroom, Waterclock*

INTRODUCTION

Why is Dwarven science always on fire?
Because normal science is boring
- User "lolghurt", Dwarf Fortress official forums

Socio-constructivist perspectives on learning suggest that learners develop domain understanding by working on authentic tasks in realistic environments. In terms of science education, this would translate as encouraging students to "act like scientists", and participate in activities of true scientific inquiry. In practice, however, creating an inquiry-based science classroom may be an unfeasible objective for some teachers. (Crawford, 1999)

As an alternative approach, research suggests that the activities of scientific inquiry can be effectively simulated and performed by students in digital environments (van Joolingen, de Jong, Lazonder, Savelsbergh, & Manlove, 2005). Current scholarship also supports the idea that digital games can be valuable environments around which students can construct both science expertise and participate in a community of practice (Barab et al., 2009). However, recent

DOI: 10.4018/IJGCMS.2015040103

reviews (authors, 2013) have concluded that digital games are rarely designed or intended to help students learn about the authentic practice of science. Shaffer (2006) proposed a form of digital game that supports "participation in a thickly authentic simulation that gives learners access to the epistemic frame of a community of practice" (p. 2). In this paper, I will describe the epistemic activities of the community of players of the game *Dwarf Fortress*, as they engage in systematic, evidence-based inquiry. I will argue that these activities are fundamentally shaped by the designed affordances of *Dwarf Fortress,* and they share the epistemic frame of practicing scientists. I will also compare their activity to the form of inquiry that is considered desirable in science classrooms.

GOALS AND PLAN OF INQUIRY

The goal of this paper is to explore and describe the emerging epistemic processes of the community of players of the game *Dwarf Fortress*. I believe this community to be unique in terms of the depth and intensity of its epistemic activity. Understanding how this community has organized its practices of knowledge generation (and how these practices differ from those found in more mainstream gaming communities) may provide a useful model for creating classroom communities around digital games for science learning. Furthermore, through analysis of how the specific traits of the game *Dwarf Fortress* mold the epistemic activities of its community of players, we may find promising design features that could be integrated into games for science learning.

To meet these goals, I will take a multi-pronged approach. First, I will provide a description of the context of inquiry, including the exact nature of the game they aim at understanding, and the distributed context in which the resulting knowledge is presented and exchanged. I will present and analyze textual artifacts of these epistemic processes, with special attention to the way in which new knowledge is expressed, reviewed, framed and reframed by a distributed, heterogeneous community. I will also conduct an extended analysis of a unique epistemic frame that some members of the community participate in, *"dwarf science",* and trace its origins and evolution. Finally, I will explain the characteristics of the game that are driving the epistemic practices of the community, and particularly "dwarf science", and what, if any, general design principles may be abstracted from this case that may be useful in creating learning environments that stimulate similar epistemic activities.

CONTEXT

About Dwarf Fortress

Dwarf Fortress (hereafter abbreviated as "DF") is a construction and management simulation (CMS) game that is freely available on the Internet. As with most CMS games, it challenges a player to build a running system, and then manage the resources it produces and consumes. Perhaps the best-known example of the CMS genre is the popular game *SimCity*, in which the system being built is a city, and the resources are money and people. The typical CMS does not have a victory condition, and so cannot be "won"; rather, the emphasis is on growth, and the player must manage their created resource economy in a way that maximizes their creative power. (Rollings & Adams, 2003). The player in a CMS game plays alone; there typically is no competition for resources or tensions created by the intersecting goals of two or more players. Because of this, the CMS game player has more space to exercise agency and creativity, beyond the basic efficiencies needed to maintain his or her creation.

In the case of *Dwarf Fortress*, a player must design and build an underground city to house a population of "dwarves", and grow an initial exploration party of seven unskilled dwarves into a thriving community with hundreds of skilled workers, warriors and nobles. As a settlement grows, it attracts more dwarves, who provide labor so that the industry can function but also drive a demand for products of ever-increasing complexity. As in all CMS games, the player does not control the dwarves directly; rather he or she gives general instructions and the dwarves obey them according to their skills and disposition to work. Realizing a player's intended actions is not straightforward: dwarves are greedy, cantankerous, despotic, covetous, lazy, alcoholic and occasionally psychotic. A skillful player must learn how to keep his or her dwarves happy, which makes them fit for work in order to produce the very things that make them happy: fine crafts and artwork, food and strong drink, valuable items and precious gems. The wealth of a settlement inevitably attracts enemies; thus, the player must direct the dwarves to construct defenses, and to continually improve them as the enemies' strength increases (hence the "fortress" part of the title). Beyond the goals of survival and wealth, the player is free to create and grow the fortress as they see fit; the only real constraints are the natural resources available to the dwarves, and the ever-increasing might of the enemy attacks the fortress's wealth attracts.

What makes *Dwarf Fortress* different among CMS games, and digital games more generally, is the depth and detail of its modeled world, whose principles operate with a regularity and complexity far beyond what most games aspire to. The game includes systems to simulate basic economic activities like farming, fishing, hunting, and a broad variety of crafts, such as smithing, masonry and brewing. Each of these activities is supported by simulations of resource growth and propagation (i.e., seeds that grow into plants that bear fruit, fish and wild game that reproduce, and predators that compete with the dwarves for the same food resources), and the behavior and interactions of materials (e.g. wood burns, iron melts, bones decompose, water flows, and when it flows into magma, it produces obsidian and steam). Geology and mineralogy are correctly simulated; metallurgy is also modeled, so that dwarves can mine tetrahedrite for copper and silver, or smelt hematite to get pig iron, and combine it with a flux stone and charcoal to make steel. Physical and mechanical systems are also simulated, so that dwarves can, for example, build windmills or waterwheels to drive millstones, or to power water pumps that fill cisterns equipped with pressure-sensitive plates that activate water-tight floodgates. The dwarves themselves are the focus of several models, such as skill, social and physiological models. Dwarves that practice a skill will grow better at it; dwarves will make friends and enemies, marry and have children, hold grudges, and be happy or sad or angry according to their individual preferences and needs. Dwarves can get sick or be injured, either by misfortune or in battle, and they heal when attended to, recover on their own, or else deteriorate and die.

Learning to play DF is not a trivial matter; the game itself contains almost no resources to aid the beginner. However, as daunting as it appears initially, a new DF player may find the detail and openness of the simulation compelling. Thus, given the paucity of included learning resources and the game's strong appeal, it is not surprising that an online community has coalesced around DF.

Communities of Players and Gaming Knowledge

In today's networked world, it is rare for any game not to attract a community of players, and one of the primary functions of these communities is to exchange knowledge specific to a given game that has value in improving the game experience (Gee, 2007). There is a certain pattern in this exchange of knowledge that applies to almost all such communities, one that reflects the purpose and value of the generated knowledge.

The vast majority of gaming communities, regardless of relative size, share knowledge in a way that is highly commoditized; skilled players produce checklists, guides, and "walk-throughs" that less-skilled players consult for the direct purpose of improving the proficiency level of their own play. The website GameFAQs.com collects thousands of these documents, called FAQs (frequently asked questions); the most popular games have 20 or more FAQs, each one detailing a particular aspect of play. A more recent phenomenon is the creation of wikis, where typically a player can find not only guidance in solving an in-game obstacle, but also a more detailed exploration of the game's inner systems, its characters, setting and mythos. This medium is more appropriate for games in which players can take multiple paths, follow different narratives, and form strongly personalized game experiences. An example of this type of game is *World of Warcraft*, the most popular multiplayer online game in the world; its main wiki, wowwiki.com, contains over 90,000 pages in 24 languages, and is the second-most used wiki in the world, after Wikipedia. Although the wiki form is more flexible and collaborative, and thus it allows for more community participation, the collected knowledge is still mostly a product of expert players' observations, and its main purpose is to instruct novices. The font of authority is almost always experience: authors of FAQs write in an authoritative voice, and the sheer volume of game knowledge (often in the form of tables, statistics, and charts) collected in their documents is a form of corroboration to their authority. For example, these are the first few paragraphs on the section of assault rifles from the 135-page long FAQ for Mass Effect[1]:

Assault Rifles

Low level players will want to look for an Avenger to swap for their Lancer because of its increased damage. All rifles at this level have poor accuracy and shots-before-overheat, so players shouldn't worry much about these categories.

Mid level characters have a broader choice of assault rifles. The Thunder does the most damage, though this can reasonably be traded for the higher SBO and accuracy of the Raptor or Tsunami. Players should find a rifle that fits their play-style and stick with it.

High level characters have the greatest choice in assault rifles. The HMWAR should be the weapon of choice for those who can afford it, but the Geth Pulse Rifle is also a fine choice--if you can find one--due to its superior accuracy. The Kovalyov is also a good choice for dealing high damage.

Ariake Technologies I II III IV V VI VII VIII IX X
DMG Tsunami - - - 216 228 240 252 264 276 288
SBO - - - 46 49 51 54 57 60 64
ACC - - - 24 26 28 30 32 34 36

The Epistemic Community of Players of Dwarf Fortress

The community of players of Dwarf Fortress, about 27,000 strong, congregates around an internet forum (www.bay12forums.com), which represents the main marketplace where knowledge of Dwarf Fortress is exchanged. In this forum, players post questions, solicit feedback and draw on each other's understanding of the game. As in most Internet forums, participants are mostly anonymous, dialogue is asynchronous in time, and older exchanges, regardless of their relevance or value, are pushed from the foreground and become more and more invisible and inaccessible with the passage of time, to be found only through deliberate searching. This logic of participa-

tion by communities of players, while not extensively studied, has been found to be place where players frequently engage in evidence-based reasoning (Steinkuehler & Duncan, 2008).

Commoditized forms of knowledge production and inscription are also present in the *Dwarf Fortress* community; after all, neophyte players benefit from this type of instruction as much as players of any other game. However, the DF community has generated a unique form. In this form, players do not describe superior forms of play, but rather investigate the modeled world in a systematic way, and present their findings for review. Authority is not derived from understanding of the game self-evidenced through the writing of a walkthrough or FAQ, but from purposeful and clever experimentation. The community of players has named this activity "dwarf science", or more commonly, !!SCIENCE!![2]. It is this particular epistemic form, arrived at spontaneously by the community and undergoing constant revision and refinement, that we center our inquiry.

RESEARCH QUESTIONS

What exactly is "dwarf science"? How is it practiced and presented? What epistemic assumptions and processes does it encapsulate? What problems does it solve for the community? And most importantly, why is !!SCIENCE!! practiced by the community of players of DF and not by players of other games, even in the case of the much more numerous, diverse and invested communities around the most popular games, such as *World of Warcraft*? What can be learned from DF and designed into learning environments, so as to stimulate learners to produce their own brand of "dwarf science"?

"Dwarf science" is a unique ontology that stands as a product of a community playing what might fairly be described as an epistemic game (Shaffer, 2006). Furthermore, "dwarf science" exists in juxtaposition not only with the commoditized inquiry present in most gaming communities, but also with the everyday business of science classrooms, and thus raises a number of intriguing questions. I will focus on two questions which highlight areas of greatest contrast between between "dwarf science" and the normal practices of science education:

1. The practice of "dwarf science" is organizing along themes that are similar to the practice of science. Why? Is it because of the adherence to a quasi-scientific structure of practice OR because players who are trained as scientists are importing real-world scientific patterns of thought and action into the greater community?
2. What motivates the community to expand the base of applicable knowledge rather than focusing on optimizing play? What is it about DF as a game that enables and sustains dwarf science?

METHODS

My research is based on observation and systematic analysis of online behavior and participation by the community of players of *Dwarf Fortress*, within the various online environments that support their efforts at scientific inquiry. Although the study of phenomena of participation in online communities is somewhat novel, there are several examples in the literature of approaches focusing on content analysis (e.g., Marra, Moore, & Klimczak, 2004) and discourse analysis (e.g., Steinkuehler & Duncan, 2008) of online forums. Virtually no empirical research has been performed featuring analysis of participation within wikis in learning contexts (e.g., Bruns & Humphreys, 2005). Despite the relative lack of methodological exemplars, This study

will use classic ethnographic methods of analytical observation to describe the rise of "dwarf science" and address the question the motivation behind "dwarf science". To address the first question, on the epistemological basis of dwarf science, I organize the evidence in the virtual space, by mapping the ecology of knowledge and tracing the paths followed by knowledge and knowledge-creators; my aim is to use this mapping to investigate the origins of the observed epistemology of dwarf science.

The present analysis is based on data collected from several streams:

- Forum discussions, as captured in the official *Dwarf Fortress* site http://www.bay12forums.com/smf/
- Wiki articles, from http://df.magmawiki.com/
- Shared artifacts of play, as present in the DF Map Archive at http://mkv25.net/dfma/

Most of the analytical effort was directed towards the DF Forums. Using the built-in search tool, I selected 10 discussion threads to serve as a sample. These threads were selected randomly from a larger search that yielded nearly 200 threads that made explicit reference to "dwarf science" and its synonyms: SCIENCE and !!SCIENCE!!. The sample yielded 476 posts by 197 different authors.

I used a database of website usage aggregate statistics, Alexa.com, to place the Forums, DF Wiki and Map Archive in their proper dimensions within the knowledge ecology, and to trace the transitions of users within this ecology, e.g. to investigate how much, how often, and from where a given website is being accessed. Since Alexa.com does not reveal the exact mechanisms with which they extract their analytics (including the sampling method), the quality of this data must be considered suspect; however, I judge it sufficiently accurate to reinforce conjectures that are warranted by other sources of data.

Analysis of Written Reports of "Dwarf Science"

In an attempt to give a more detailed description of the how the DF community investigates DF, I will analyze cases of reports of "dwarf science", and attempt to draw from these reports the essential elements towards understanding the epistemic processes at play. As a working frame of reference, I will use methods of descriptive analysis of text featured in the work of Charles Bazerman. In *Shaping Written Knowledge*, Bazerman (1988) describes the evolution of the experimental report as an account of scientific activity. In the chapter "Reporting the Experiment", Bazerman reviews 100 articles from the *Philosophic Transactions of the Royal Society of London* spanning more than 150 years, and provides an analysis based on several guiding questions; the ones that are most relevant to the present analysis may be paraphrased as:

1. To what kind of "event" does the experimental report refer? How is the term "experiment" used?
2. How fully and in what manner are experimental events described?
3. How fully is the methodology described? How fully and in what way are methodological concerns discussed?
4. How precisely and completely are results presented? How much and what kind of discussion and interpretation are present, either in the main report or in any replies posted on the forum by others?
5. How is the account of the experiment organized?

The historical lens of Bazerman is appropriate for analyzing reports of "dwarf science" for two reasons. First, if the evidence available suggests that the community is not so much importing the template of an experimental report, but rather refining it from earlier forms and modifying it to suit different rhetorical, epistemological and sociocultural purposes, then this dynamic has already been described by Bazerman as having played out over 150 years of *Transactions* reports. Also, if the community of players of Dwarf Fortress has generated "dwarf science" as a method of querying their reality, then its reports are reflections of the goals and problems of that method. If, as Bazerman suggests, the genre of experimental reporting evolved to solve problems inherent to experimental science (Bazerman, 1988, p.63), and the epistemic characteristics of the latter can be inferred from analysis of the former, then (with the necessary modifications) we may make claims about the nature of "dwarf science" by analyzing its reports in a similar fashion.

FINDINGS

Dwarf Science as a Response to Systemic Uncertainty

The complexity of the DF modeled world and the depth of its simulation means that the learning curve for the game is uncommonly steep. Even the most basic game challenges encountered during the first few hours of play require extensive scaffolding; furthermore, there are specific in-game events that, when game conditions are met, trigger an explosion of systemic complexity. For example, when the fortress wealth reaches a certain number, it will attract a hostile siege, which the player must respond to or watch their fortress succumb to invaders. Or, when the fortress reaches a certain population, the "dwarven economy" will activate, introducing issues of monetary policy, wealth disparity and allocation of labor, forcing the player to mint coinage, set tax rates and deal with the supply and demand of goods.

Failure to grasp, accommodate, and plan for these moments of complexity carry a steep penalty in game terms. Often a player has no recourse but to watch their fortress collapse, flood, fall prey to rampaging beasts, and witness how the dwarfs starve, die in battle or go mad. In fact, the threat of catastrophe is so constant, and its causes so diverse and occasionally unfathomable, that one of the DF community's unofficial mottoes is: "losing is fun"[3]. Thus, to avoid the negative aspects of "fun", it becomes crucial to understand how the modeled world works, and the principles under which it operates, so that play of the game may advance.

So a Tundra Titan attacks my fort... I send out my military, and they fell it, losing only 3 of their force of 20! Victory and celebrations right??? Nope.
They return to their giant golden (literally) feast hall, with adamantine statue and all, and bury the dead in luxurious crypts with godly coffins... And proceed to fist-fight and uproot everything and everyone, downward spiral of death and depression that wipes my 80+ dwarf civilization. What did I do wrong? What possibly could have turned this situation around?" – talrave, 12/3/2010[4] (emphasis added)

As sophisticated as the world of *Dwarf Fortress* is, it is still a simulation and must perforce have limits on the scope and detail of the phenomena it represents. These approximations and issues of scale and grain size act together to produce what I have termed *systemic uncertainty*: players can never be sure of exactly *how exactly* and *how realistically* the modeled world operates, only that it *does,* to a variable yet knowable extent. In most games, the rules of operation of the modeled world are introduced gradually, with options for play appearing only as they

become necessary to solve the specific game challenges present at a given time, in an upwardly increasing level of complexity. However, there is a manageable limit to this complexity, so that most digital games can (for the most part) be solved in the mathematical sense: the entire game-space mapped, an optimal solution found, and prescription for the solution made available. This process of solution by expert players is a common feature of computer games, and forms the central content of most websites and magazines devoted to gaming. No such solutions exist for DF, however. The level of systemic uncertainty present in DF makes it nearly impossible to predict, or prescribe, solutions to all possible game challenges.

The Origins of Dwarf Science

Given that optimal play of DF is not really possible, and "fun" (losing) probably not avoidable, the community places less value on producing documents that encode optimal-play instructions, such as those found on GameFAQs.com. Instead, an earlier community trend was for more skilled players to focus on building "megaprojects" that showcase their ability. These "megaprojects" often revolved around intricate designs of dubious function or use; they later became known as "stupid dwarf tricks", defined in the DF Wiki as "any project that requires a large amount time and effort - often for little or no practical benefit. They exist only as a challenge for experienced players."[5] While impressive, the DF Wiki's list of "stupid dwarf tricks" highlighted two facts: first, most of the "tricks" were unverified opinions of what *should* work, or accounts of potentially fictitious accomplishments; second, that it was far from understood *why* "stupid dwarf tricks" worked, as many of them depended on game phenomena that operated strangely, unpredictably, or counter intuitively.

This empty space was filled with "dwarf science", which has its origins as a simple report by a player of the results of their experimentation on a particular "stupid dwarf trick"[6]. This discussion thread, titled "I committed SCIENCE!", and subsequent threads, had the catalyzing effect of recasting the terms in which game knowledge was being shared, drawing discussion away from the "how" of "stupid dwarf tricks" into the experimental, evidence-centered "why".

Science is about discovering how a mechanic works, either it is how far do kittens fly when on a bridgeaoult [a fast-rising drawbridge that functions as a catapult, "bridgeapult"], or from what height do kittens splat horribly wonderfully. All experiments should be documentated through a log. Stupid dwarven tricks are the practical application of the discoveries from Science. If you want to build a machine that sends thermonuclear cats at your enemy, you do not want to scroll through pages of algorythm of experiments on how much cats do smoke, or how many tries have been done to find out that. You juts want to smash someone's head with a cat with exploding uranium bones! The difference is pretty clear: just like difference between theoretical physics and experimantal physics." – wagawaga, 11/27/2009[7]

This thread also contains an initial play of definitions, as the community attempts to define what "dwarf science" should look like.

[dwarf science] must be observable. You must be able to SEE genocide/death/pain/suffering/magma-related-pain occurring.
It must be repeatable. You must not be satisfied with only one dead/injured kitten/goblin/elf/species, you must repeat the experiment.
It must be emperical. Experiments must be done, or at the very least, a design for a fort to perform this experiment must be created.

It must be verifiable. Many other fort leaders must be able to repeat your results... in fact, they are encouraged to!!" – digitCruncher, 11/26/2009[8]

Now we know what the community set out to do, what it initially identified as the problems that needed to be solved, and how it envisioned the desired characteristics of their epistemic work. What follows is to examine the specific form of the products of the epistemic work, namely written reports of "dwarf science", submitted for public scrutiny to the DF official forums.

Analysis of Dwarf Science Reports

This section contains in-depth analysis of two cases of "dwarf science" reports. These reports were chosen for their representative nature; Case 1 standing for the origin point of dwarf science (a "stupid dwarf trick"), and Case 2 being more representative of dwarf science in its currently realized form.

Case 1: The "Water Clock" Case

The first case for analysis will be a report authored by the user "Dood" on July 2nd 2010, titled "The dwarfy way to build a waterclock"[9]. In this report, the player presents an account on how a certain configuration of states can, when modified slightly, produce an event that happens at regular intervals, i.e. a "clock".

I think I just managed to build a new kind of waterclock.
[...] The first layer is completely filled with water, as are big parts of the second one. The third one contains more dry soil and a conveniently placed magma pipe. What intrigued me was the fact that directly above the magma pipe the caverns were only separated by floor tiles, with no "solid" rock in between. So I hatched a plan: Drain layers 1 and 2 directly into the magma pipe. A quick cave-in later the dwarven waterclock was born. (See Illustration)
What I thought what would happen (all the water from the first two cavern layers rushing down and completely filling the third cavern) didn't. Instead I have a slow but steady stream of 7/7 tiles of water dripping into the the center of the magma pipe. Each 7/7 hitting the magme (sic) creates a huge plume of steam and a single tile of obsidian, which immediately drops down to the bottom of the magma sea and somehow vanishes. Therefore the pipe is not being capped by the usual crust of obsidian, which means the clock will keep ticking until the upper caverns run out of water.
Well there you have it: Precision timing using all of the four classical elements plus SCIENCE!

Analysis of the "Water Clock" Case

The event described in this report is not so much an experiment as it is a case of the simulated systems interacting to produce a result that the player did not expect. The author has a hypothesis of how the magma and water will interact, but instead of getting either of the expected outcomes (the water "rushing down" and the "usual crust of obsidian" capping the magma pipe), a phenomenon is presented that the author quickly frames as a novel application (i.e., a "water clock").

Bazerman notes that, in the first volume of *Transactions*, "a number of experiments reported are simply cookbook recipes for creating marvelous effects or effects of practical use." (p. 66) This description matches the described case: the author of the report intends to present an interesting interaction with mostly novelty value. Accounts of these kinds of novel applications

Figure 1. A quick cave-in later the dwarven waterclock

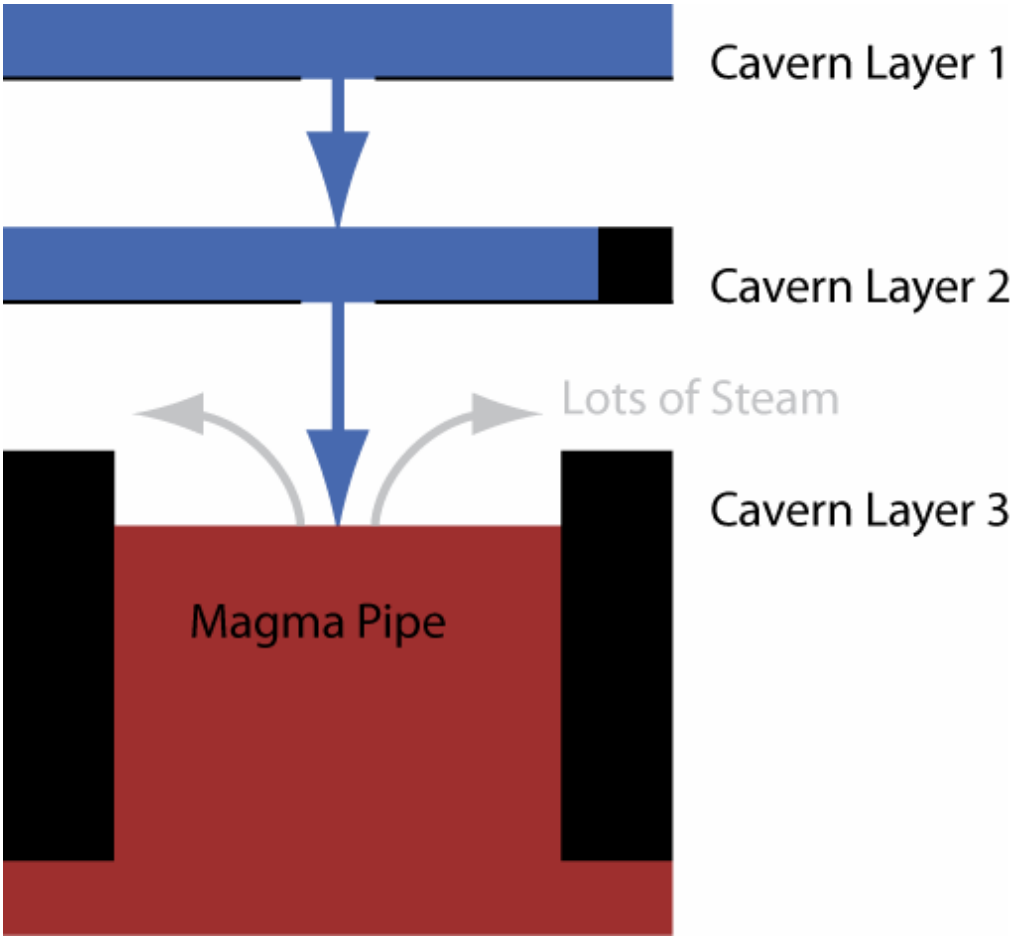

are somewhat common in the DF community, where they are called "stupid dwarf tricks", i.e. amusing phenomenon which can be applied to situations of no practical value. The DF Wiki lists over 50 examples of "stupid dwarf tricks"[10], noting that these applications serve only as challenges for experienced players.

In terms of method, the author of the report is somewhat vague. He or she describes an existing situation (two z-levels nearly filled with water, with a magma pipe beneath); furthermore, the action of the experimenter in bringing about the result is extremely brief ("a quick cave-in later"), and it is in fact not explicitly stated whether the player caused the cave-in. We can find parallels to this terseness in the early volumes of the *Transactions*, in which, according to Bazerman, how an experiment was performed is not generally detailed, or mentioned only in passing. The "water clock" author clearly expects the readers to know how to cause a cave-in, or perhaps does not believe that the method of creating the cave-in would have any bearing on the result. In any case, it is the natural state of the system (layers of water over the magma pipe), not the cave-in, that is the means of producing the phenomenon of interest. The "water clock" account is, in fact, organized in such a way that the players hand is de-emphasized. After describing the

initial state, the author inserts a diagram, then a narrative account of the phenomenon that was produced. The diagram serves to focus the reader's attention on the undisturbed state; the sections on the cavern floors that the player opened are nearly invisible. Arrows serve to visually reinforce the author's narration of the unexpected result of this experiment. Only a single trial is presented; the narrative is linear and jumps quickly to the conclusion that this phenomenon will feature "precision timing", when in fact, no evidence of this timing is presented, only that the magma-to-sinking obsidian transformation "will keep ticking."

All told, this account fits well with Bazerman's description of early early-volume experimental reports from *Transactions*. In those volumes, a scientific report was "simply a matter of news" (p. 77), and not a matter of contention. Likewise, the "water clock" account is not challenged in the forum replies; of the 10 responses posted to the account, none of them question the results, only to note that similar situations (e.g. a reservoir or an underground river over the magma pipe) would likely produce the same result.

Case 2: The "Lignite Block in an Iron Bin" Case (a.k.a. the "Reactor")

This case was selected because it presents a number of interesting contrasts to the "water clock". In this account, dated February 20th 2009, the forum user "Kanddak" recounts their efforts at understanding the unexpected behavior of a block of lignite locked in a fire-proof box[11]:

Summary: Lignite blocks put in an iron bin and submerged in magma ignite and can't be extinguished, bin thenceforth sets fire to dwarves and instantly vaporizes water.
So I heard that you could turn blocks back into stones by putting them in a magma-safe bin; submerging said bin in magma, thereby melting the blocks; then recovering the bin and its contents and allowing the molten blocks to cool into stones.
I set out to test this today. It didn't seem to work. It turns out there's a difference between a "dolomite", which shows "Uses: Make steel bars, Make pig iron bars" and is a stone, and a "Dolomite", which has no uses and shows up in "globs" on the stock menu.
When I tried this, I put two Lignite blocks in my bin with some other stuff. Rather than melting, they caught on fire inside the iron bin and began emitting smoke.
As soon as I opened the test chamber to try to recover the bin and see whether I could persuade my smelter to take a big-S Sphalerite, the first dwarf to set foot in the chamber caught fire and died. I poured buckets on it for a while, but that didn't seem to change anything. The bin shielded the burning blocks from being extinguished by water.
I built a system to pump water into the chamber to see what would happen.
It instantly drained my reservoir, so I activated my fill pump from the cave river and started pushing more water in. The bin of doom is still instantly vaporizing all the water I pump towards it from a tile away. (Bin tile: Dry. Next tile: Dry. Second tile: Water [6/7].)
I just pitted a rhesus macaque child onto it, and it died in the heat, then quickly became bones.
I can't move the thing, but I could sure have some fun with it.
I'd like to see what happens if you make one on a retracting bridge and then drop it into a lake, or if you put it in a 1-wide passage that's the only entry to your fort during a siege or elf caravan.
It could potentially have magma's advantages of burning goblins to death and destroying their narrow clothes, without the trouble of draining the magma when you want to collect the iron stuff, due to the flaming goblins running past the bin before dying.

In this case, the player has similar goals at the outset as the player in the "water clock" account: both discover an interaction that does not work as expected, and in both accounts,

we are offered an explanation. For "Kanddak", the reason the first action fails to produce the expected results is that the game treats "dolomite" and "Dolomite" as two different substances. This is perhaps a useful finding, but the author is more interested in reproducing another fact, that blocks can be transformed back into stones by putting the blocks in the bin and covering the bin in magma, thereby melting the block and returning it to its natural state when it has cooled.

Both authors make very clear where exactly their observation diverged from what they expected. As with the obsidian forming, then sinking, in the "water clock" case, the author of this report notes that lignite did not melt, but caught fire, and presents additional evidence towards that fact ("and began emitting smoke"). In following through with the experiment as originally intended, the author finds another unexpected result: the bin is hot enough to burn the dwarf that entered the room to try to recover the melted minerals.

At this point, the author begins what we can more confidently label as a series of experimental trials, in which the player is positioned as the author of purposeful actions ("I built"… "I activated"… "I pump", and so on). The report contains four different experiments, all aimed at uncovering the properties of this mysteriously burning bin. The result of at least one of these experiments is stated precisely: the bin is hot enough to vaporize any amount of water 1 tile away.

The way in which these experiments are carried out, and the purpose behind them, matches Bazerman's descriptions of the reports in Volumes 5 and 10 of *Transactions*. Bazerman characterizes these accounts as "conscious investigation of phenomena involving some doings or manipulations" (p. 66), but also notes that "nowhere do the reports… suggest that questions, theories, problems or hypothesis are being explicitly explored or tested." This last clause might be seen as less applicable, until we realize that "Kandakk" is not investigating *why* the lignite locked in the bin burns and cannot be extinguished, nor is he or she advancing any hypothesis as to the cause, instead offering possible applications.

There is far more controversy within the DF community regarding the "reactor" case than the "water clock", and the nature of the questions and replies are very different. The author of the report is asked for further observation ("how long does it burn?") and expansion of the experimental scope ("does this work with wooden blocks as well?"). In response to these questions, "Kandakk" reports:

When my first bin quits, I'm going to do some more conclusive tests on fires, by filling my pit with each of the following materials, first by themselves and then in a bin:
Logs.
Wooden blocks.
Charcoal.
Lignite.
Lignite blocks. (outside the bin only, since I already know what it does in the bin)
Coke.
Green glass blocks. (as a control)
I'll also put a hatch in the bottom of the pit first. Then I'll drop the flaming stuff down a z-level, pile more stuff on the hatch, and drop it in to confirm that dropping in another piece of coal should be enough to keep a fire going.
Maybe I'll try booze in an iron barrel, too, or experiment with some more inventive ways to try to put water on it and see if there are some ways to put out fires that work better than others (i.e. at all, in the case of burning coal in a bin), or significant difference between flaming items inside and out of a bin.

Also, a second user takes up this finding and begins to build on it, reporting his or her progress to the community. User "Akhier the Dragon hearted" expands the inquiry by trying to use the "reactor" to evaporate an entire ocean (and succeeds); in fact, it is this user who provides most of the physical evidence in favor of the "reactor" in the form of screenshots, animations and diagrams, although he or she offers no analysis to go with the new evidence. Several other users also reported variants of the "reactor" or novel applications for it; in total, the original posting generated over 200 responses.

The "reactor" account is organized in a very different way than the "water clock" case. In the latter, the author is simply reporting a phenomenon: When B is happens to A, X happens. In the case of the "reactor", the structure is much less linear; although the "A plus B results in X" format is still present, the author wants to convey that the phenomenon X was understood through several conscious manipulations, and (particularly in the follow-up reports) that efforts are being taken to take variables into account.

Epistemic Spaces of Dwarf Science

An important question is whether the tools being used to discuss, refine and disseminate "dwarf science" are as important as the systemically-complex environment that originates its observations and insights. To address this question, I will describe the way in which the DF community uses its communication tools, and how these tools interact in a kind of "knowledge ecology".

The *Dwarf Fortress* community of players mainly accesses three external sources of information when playing the game: the official forums, the DF Wiki, and the Map Archive. Each of these elements is described in detail in individual sections.

The Official Forums

The creators of *Dwarf Fortress* operate and maintain an online forum for players to discuss issues related to their experiences playing the game. The forum has 22,000 registered users and is fairly active; since 12/1/2007, the forum has more than 1,782,000 posts across more than 72,000 topics. The DF community can be said to be highly active in the forum compared to other communities dedicated to more "mainstream" games, which typically sport far large user bases: for purposes of comparison, the official forums for the top-selling game *Rock Band* have 462,000 members yet only 3.8 million posts. Forums allow a range of modes of participation that are largely modulated by context, topic and purpose of communication (Thomas, 2002). And since player activity on the forums is rarely confined to one purpose or mode, most players in the sample expressed opinions over a wide range of topics.

The DF Wiki

The DF Wiki, like all wikis, is completely open to editing, so that any user can edit any page, adding or removing content at will. While this sounds like a recipe for knowledge destruction, one important feature is the ability of the community to assign a "quality rating" to each article, which functions as a mechanism of peer-review. DF Wiki articles are rated by readers along a four-point scale (Tattered, Fine, Exceptional, Masterwork) according to several criteria:

- Comprehensively covers an important topic
- Contains no unverified information
- Has an appropriate number of outbound functional links
- Has multiple editors

- Is properly categorized and aesthetically pleasing.[12]

These criteria operate in concert to drive the DF Wiki towards *veracity* and *orderly storage of information*. In requiring that each article have multiple editors and that information be properly verified, the DF Wiki follows a template from experimental science publication; the criteria of comprehensive coverage, linking and categorization corresponds to a template from scientific book publishing, particularly of handbooks and textbooks. As of this writing, the Wiki contains 1181 articles, of which 55.5% are rated Exceptional or above.

As valuable a resource as the DF Wiki has become, it is interesting to note that it was not always so. It is common for players to be skeptical of the contents of the DF Wiki related to dwarf science, much like research scientists tend to view school textbooks,

The wiki is notoriously inaccurate on subjects at the cutting edge, frequently reflecting passing memes, folklore, or the word on the street instead of true dwarven science. –zchris13, 11/11/2009.[13]

The Map Archive

The least prominent, but nonetheless important, epistemic space used by the DF community is the Map Archive, a website dedicated to storing and sharing games in-progress. In the Map Archive, players can upload and share working examples of complete fortresses or smaller subsystems, e.g. defensive fortifications, factories, dormitories, etc. Unlike the DF Wiki and the Forums, the Map Archive is not a discursive space. It is used almost exclusively as a repository for artifacts used in discussions within the Forums, not as demonstrations by articles in the Wiki. This leaves the Map Archive as smaller and less dynamic of the three spaces in terms of epistemic activity. However, it fills the community's need for a readily-available set of exemplars that can be referenced in a consistent and user-friendly way.

DISCUSSION

More on the Analyzed Cases of "Dwarf Science"

I have tried to draw two cases of accounts of "dwarf science" that are representative and that can be productively analyzed; however, there are over 200 threads on the DF forums that contain reports of "dwarf science", along with several thousand replies, and so capturing the diversity and depth of epistemic activity of the community must remain outside the scope of this work. However, the salient features of these cases, and their close kinship with Bazerman's characterization of early scientific reports in *Transactions*, make it possible to extend some of Bazerman's arguments to cover the epistemic processes of the DF community.

Bazerman states that "science will as well stand up for scientists, for the authors also subordinate themselves to scientific knowledge as currently constituted" (p. 30), and this seems no less true for "dwarf science" than for "real" science. Players in both the examined cases are intent on holding up their achieved knowledge of *Dwarf Fortress* to the standards of "dwarf science"; authors of FAQs and walkthroughs rarely make their expertise accountable to the community in this way, nor are they necessarily concerned with how the particular products of their expertise fit into the existing body of knowledge. In contrast, DF players who write about !!SCIENCE!! are singularly concerned with positioning their contributions within the known field of DF knowledge.

That said, an important divergence between the DF accounts and Bazerman's historical perspective on the rhetorical functions of scientific writing is that, unlike the *Transactions*, the evolution of "dwarf science" report writing does seem not follow a chronological path. The "water clock" case, for example, has more in common with an early-volume *Transactions* report, even though it came almost a year *after* the "reactor" report, which is more akin to a later-volume *Transactions* paper. This divergence is not fatal to the analysis, however, when we consider two facts. First, the DF community contains a broad mixture of novice and expert players, with varying positions of centrality within the community (cf. Lave & Wenger, 1991). User "Kanddak", for example, can be said to me a much more central figure in the DF community, as evidenced by their 649 posts created over 2 and a half years; in contrast, user "Dood" has only 15 posts. These different levels of participation may account for an increased commitment to what, in the eyes of the community, does and does not constitute "dwarf science". In other words, it may be the case that Bazerman's chronological evolution happens at the level of the *player* (as he or she is socialized to the assumptions and expectations of "dwarf science") rather than at the level of a community of practitioners, whose written artifacts Bazerman originally studied.

Follow-Up: Is Dwarf Science Modeled After Real-World Science?

A possible explanation for the emergence of dwarf science is that, to address the issues of systemic uncertainty, the community is importing templates of 'real-world' science whole-cloth. This would be the case, for example, if a portion of the DF community had some experience in the everyday practice of science and, unsatisfied with the haphazard production and dissemination of knowledge present as "stupid dwarf tricks", began to drive more evidence-centered, structured inquiry, i.e., dwarf science.

This alternative explanation for the origin of dwarf science, while not contrary to the main finding, would undermine the value of "dwarf science" as a model for classroom inquiry around epistemic games. "Dwarf science", in spite of its unique setting, is simply an artifact produced by players who happen to be scientists applying to DF the discursive and epistemic forms they already know works for science. Yet if this scenario were to be true, the proportion of forum users who fit the demographic of scientists (i.e. older than 30 years of age and holding graduate degrees) would closely approximate the proportion of the DF community that is engaged in dwarf science; this would be taken as strong evidence that the practice of inquiry originates, and is largely bound, in this section of the population who are already familiar with the practice of science.

It is virtually impossible to determine the demographic composition of only the subset of DF players who engage in "dwarf science". But the aggregate data for all users of the forums is available (Figure 2). The demographic data indicate that the DF community, while slightly older, is not substantively different in composition from a freshman college science class in terms of academic achievement, with graduate degree holders over the age of 35 being greatly underrepresented. However, this doesn't necessarily disprove the alternative hypothesis that a very small fraction of players engaged in dwarf science are drawing on practices and templates that originate in their professional lives; this demographic could still be present in the population, accounting for the majority of the activity of inquiry. Therefore, the key to addressing this question is: what proportion of the DF community is actively engaged in dwarf science?

Answering this question is far from straightforward, but the data allows us to at least make a conjecture on the basis of facts. In a random subsample of 5 discussion threads (taken from our original sample of 10) dedicated to dwarf science, I found the following:

Figure 2. Demographic data of users of the DF Forums. Bars indicate over- or under- representation of that class of DF official forum visitors relative to the general population. Source: Alexa.com

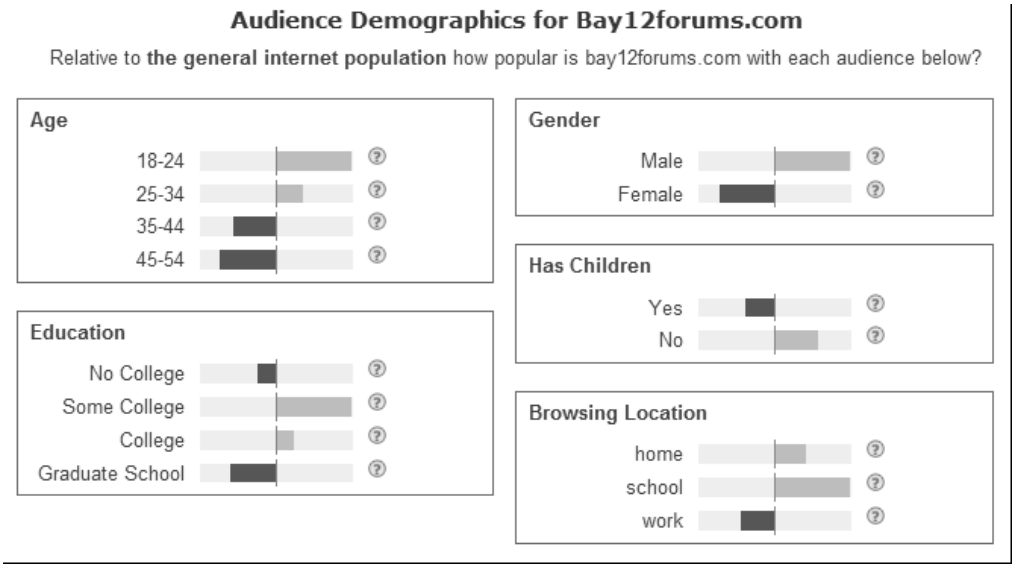

From search results, we know there are around 200 discussion threads devoted to dwarf science. From the sample data above, we know that there is not a very large overlap between threads in terms of participants; players participate in the threads they find interesting or in which they feel they can contribute. If the ratio of unique users per discussion thread is conserved, then we can estimate that around 5000 registered forum users (23%) so far have participated in discussion threads related to dwarf science. Such a large percentage is suggestive, if not definitive, evidence that the DF community's inquiry efforts have a broad base of participation, not limited to a small group of a given demographic.

It may be argued, however, that the registered DF Official Forum users are not representative of the greater community of DF players, with Forum users possibly being more dedicated and active in the community than those who simply browse the Forums anonymously. Participation in "dwarf science" would then select for players who are dedicated versus those who are more

Table 1. Authorship and impact data of the sample subset

Thread	URL	Unique Users	Comments	Times Read
1	http://www.bay12forums.com/smf/index.php?topic=67504.0	28	50	2198
2	http://www.bay12forums.com/smf/index.php?topic=46033.0	41	84	13298
3	http://www.bay12forums.com/smf/index.php?topic=51627.0	27	62	1980
4	http://www.bay12forums.com/smf/index.php?topic=42111.0	20	46	3125
5	http://www.bay12forums.com/smf/index.php?topic=42721.0	18	35	2960
	TOTAL	127	277	23561

casual. This objection is fair, however, it should be noted that there are indicators (mostly in the form in user-created polls) that indicate that the population of DF Forum users is in fact *younger* than the Alexa sample[14], making the gap between the observed "DF Forum user" demographic and the hypothetical "real-world scientist" demographic even larger, further weakening the case in favor of whole-cloth importation of templates of inquiry.

CONCLUSION

By the large, *Dwarf Fortress* players understand that they are playing a game, but they see in their game a detailed, opaque modeled world that operates on regular, knowable rules. Many of these sets of rules also have clear analogues in the real-world, so in many respects, the modeled world behaves in a way that is both familiar and understandable in terms of real-world science. However, some elements of the game are only approximations of real-world analogues, rather than accurate models. In other words, players are prompted to think about real world facts, and how they may be instantiated in the game, but are consistently limited in transforming those facts into improved play because of the boundaries enforced by the simulation. The "water clock" author expects an obsidian crust to form on the magma, but does not expect it to sink; the "reactor" author expects lignite to burn, but does not expect the fire to be unquenchable.

I have termed the unknown, persistent, and unresolvable slippage between the game-based functions and its real world analogues *systemic uncertainty.* It is this uncertainty that "dwarf science" seeks to resolve, and thus it deals with phenomena that, in the light of our everyday understanding of the world, are surprising, counterintuitive, bizarre, or humorous. When players discuss phenomena, it is easy to lose track of whether they are referring to real-world or game-specific events, or a conjectured similarity between both. The result is an in-between state akin to what (Latour, 1999) called a "circulating reference", only that the elements that circulate through the transformative chain are not only real-world scientific phenomena and their representations, but also the game-instantiated analogues thereof, and *their* representations.

I believe that the form of systemic uncertainty that permeates *Dwarf Fortress* is the design trait that gives form and impetus to the rich activities of inquiry that constitute "dwarf science", and that perhaps such a design trait would have a similar effect in other learning environments. This hypothesis would run counter to a current trend of digital environments for science learning toward representing phenomena in more directly observable way that reveals the inner workings of the model; i.e. a "glass-box simulation". In contrast, *Dwarf Fortress* is a "black box" simulation, in which the players access only the representations of phenomena and not the underlying model. Could something akin to "dwarf science" emerge among users of a "glass box" learning simulation? I am not so certain. What would be the impetus for such inquiry? Phenomena as apparently mysterious as the ever-burning bin in the "reactor" case could be easily explained away if DF were "glass boxed", and what would be lost would be the dynamic, open-ended activity that centers around probing an opaque, occasionally cryptic, but mostly knowable world, and the sense of adventure and risk that goes with it.

The sense of adventure inherent in trying to understanding the world through experiments that always teeter on the verge of chaos (or "fun") is encoded in the name "dwarf science", and the phraseology "for !!SCIENCE!!" as it has emerged and is currently used in the community. *Dwarf Fortress* players understand that the kind of inquiry they are performing, and that it is, in some form an instantiation of thing that exists outside the game world, called "science". However, the DF community signals a trait about "dwarf science" in the way that the term is used. Common formulations, such as "I committed !!SCIENCE!!!" or "for !!SCIENCE!!" indicate

that players do not see "dwarf science" as a static body of knowledge, but as an accountable, purposeful activity, a heroic pursuit of knowledge that maintains a mischievous and playful perspective. Science on fire, as it were.

Ultimately, "dwarf science" serves as an authentic (if idiosyncratic) example of the kind of inquiry that, according to reform perspectives on science learning (Duschl, Schweingruber, & Shouse, 2007; NRC, 2011), should be forefronted in science classrooms. Furthermore, my observations of the DF knowledge spaces support the view that these spaces roughly map onto three types of resources available to students engaged in inquiry, with the Forums standing in for "classroom discussion", the Wiki as a sort of "textbook" and the Map Archive as a "science fair." These three knowledge resources are deeply connected to the practice of "dwarf science", each fulfilling a different function, but providing epistemic resources that the other spaces largely cannot.

The dynamism and intellectual ferment present in the DF community as represented by their interest and active participation in dwarf science is not fully explained away by issues of demographics or self-selection. Rather, it is the unique set of affordances of DF as a complex modeled world, capable of containing and sustaining a level of systemic uncertainty, which drives the DF community's epistemic activities. Efforts by the DF community to resolve this uncertainty have given rise to a form of evidence-driven, peer-reviewed, community-curated store of knowledge about the modeled world that, while relevant only in within that world, can rightly be called science. Or, for that matter, dwarf science.

REFERENCES

Barab, S. A., Scott, B., Siyahhan, S., Goldstone, R., Ingram-Goble, A., Zuiker, S. J., & Warren, S. (2009). Transformational Play as a Curricular Scaffold: Using Videogames to Support Science Education. *Journal of Science Education and Technology*, *18*(4), 305–320. doi:10.1007/s10956-009-9171-5

Bazerman, C. (1988). *Shaping Written Knowledge : The Genre and Activity of the Experimental Article in Science*. University of Wisconsin Press.

Bruns, A., & Humphreys, S. (2005). Wikis in teaching and assessment: The M/Cyclopedia project. In *Proceedings of the 2005 international symposium on Wikis* (pp. 25–32). ACM. doi:10.1145/1104973.1104976

Duschl, R., Schweingruber, H., & Shouse, A. (2007). *Taking science to school: Learning and teaching science in grades K-8*. Washington, DC: National Academies Press.

Gee, J. P. (2007). What Video Games Have to Teach Us About Learning and Literacy. Second Edition: Revised and Updated Edition (2nd ed.). New York, NY: Palgrave Macmillan.

Latour, B. (1999). *Pandora's hope: essays on the reality of science studies*. Harvard University Press.

Lave, J., & Wenger, E. (1991). *Situated learning: legitimate peripheral participation*. Cambridge University Press. doi:10.1017/CBO9780511815355

Marra, R., Moore, J., & Klimczak, A. (2004). Content analysis of online discussion forums: A comparative analysis of protocols. *Educational Technology Research and Development*, *52*(2), 23–40. doi:10.1007/BF02504837

NRC. B. on S. E. D. of B. and S. S. and E. (2011). A Framework for K-12 Science Education: Practices, Crosscutting Concepts, and Core Ideas. National Research Council. Washington, DC: The National Academies Press.

Rollings, A., & Adams, E. (2003). *Andrew Rollings and Ernest Adams on Game Design (Ltd Rmst.)*. New Riders Games.

Shaffer, D. W. (2006). Epistemic frames for epistemic games. *Computers & Education*, *46*(3), 223–234. doi:10.1016/j.compedu.2005.11.003

Steinkuehler, C., & Duncan, S. (2008). Scientific habits of mind in virtual worlds. *Journal of Science Education and Technology*, *17*(6), 530–543. doi:10.1007/s10956-008-9120-8

Thomas, M. J. W. (2002). Learning within incoherent structures: The space of online discussion forums. *Journal of Computer Assisted Learning*, *18*(3), 351–366. doi:10.1046/j.0266-4909.2002.03800.x

Van Joolingen, W. R., de Jong, T., Lazonder, A. W., Savelsbergh, E. R., & Manlove, S. (2005). Co-Lab: Research and development of an online learning environment for collaborative scientific discovery learning. *Computers in Human Behavior*, *21*(4), 671–688. doi:10.1016/j.chb.2004.10.039

ENDNOTES

1 http://www.gamefaqs.com/pc/944902-mass-effect/faqs/51033
2 The orthography is derived from the in-game text representation of an object on fire (e.g. a burning tree would appear in the game as a !!TREE!!). The significance and implications of this orthography are described in detail later in the paper.
3 http://df.magmawiki.com/index.php/Fun .
4 http://www.bay12forums.com/smf/index.php?topic=72224.msg1770068#msg1770068
5 http://df.magmawiki.com/index.php/DF2010:Stupid_dwarf_trick
6 http://www.bay12forums.com/smf/index.php?topic=45268.0
7 http://www.bay12forums.com/smf/index.php?topic=45389.msg888892#msg888892
8 http://www.bay12forums.com/smf/index.php?topic=45389.msg887898#msg887898
9 This report can be found, along with the community responses to it, at http://www.bay12forums.com/smf/index.php?topic=60641.0
10 http://df.magmawiki.com/index.php/DF2010:Stupid_dwarf_trick
11 http://www.bay12forums.com/smf/index.php?topic=31244
12 http://df.magmawiki.com/index.php/Dwarf_Fortress_Wiki:Quality
13 http://www.bay12forums.com/smf/index.php?topic=44667.msg863198#msg863198
14 http://www.bay12forums.com/smf/index.php?topic=69912.0