THE DESIGN, DEVELOPMENT AND ASSESSMENT OF AN EDUCATIONAL SPORTS-ACTION VIDEO GAME: IMPLICITLY CHANGING PLAYER BEHAVIOUR

by

Chad Ciavarro
B.Sc., Simon Fraser University, 2002

THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

In the School
of
Interactive Arts and Technology

SIMON FRASER UNIVERSITY

Spring 2006

APPROVAL

Name: **Chad Ciavarro**

Degree: **Master of Science**

Title of Thesis: The design, development and assessment of an educational sports-action video game: Implicitly changing player behaviour

Examining Committee:

Chair: ______________________________

Dr. Rob Woodbury
Chair

____

Dr. Mike Dobson
Senior Supervisor

____

Dr. David Goodman
Supervisor

____

Dr. Janet McCracken
Supervisor

____

Dr. David Kaufman
External Examiner

Date Defended/Approved: Dec. 12/2005

Simon Fraser University

Ethics Approval

The author, whose name appears on the title page of this work, has obtained human research ethics approval from the Simon Fraser University Office of Research Ethics for the research described in this work, or has conducted the research as a member of a project or course approved by the Ethics Office.

A copy of the approval letter has been filed at the Theses Office of the University Library at the time of submission of this thesis or project.

The original application for ethics approval and letter of approval is filed with the Office of Research Ethics. Inquiries may be directed to that Office.

ABSTRACT

Concussion education and prevention for youth hockey players has been an issue of recent concern amongst sport medicine practitioners and hockey's administrative bodies. This thesis details the design, implementation and assessment of a sports-action hockey video game that aims to reduce the aggressive and negligent behaviours that can lead to concussions. The game, termed *Heads Up Hockey*, was designed to modify game playing behaviour by embedding an implicit teaching mechanism within the gameplay. Educational games often suffer from the problem of indirection, that is, the content the learner is intended to learn is indirectly related to the gameplay. With *Heads Up Hockey*, participants were expected to learn by simply playing to win, in contrast to playing to learn. The 21 participants in the experimental learning group significantly improved their mean score on a composite behaviour indicator ($p = 0.0002$) compared with no significant change amongst the 21 control group participants.

ACKNOWLEDGEMENTS

Near the end of my undergraduate studies, I found myself walking through the SFU Computing Science building with Dr. Mark Drew. He remarked to me, without provocation, that if I was to go to graduate school I had better not wait. As it turned out I did wait one year, but those words remained with me. In reality, I never actually left Simon Fraser University, finding gainful employment in Kinesiology courtesy of Dr. David Goodman and Dr. Stephen Robinovitch. Eventually it was my time to stop waiting, and plunge head first into graduate school.

Dr. Mike Dobson was my first contact at SIAT and helped me to continually revise my application until it was satisfactory. Over two years later, not a lot has changed and Mike is still asking for revisions. In the same vein, Dr. Goodman has provided much needed direction and support for my graduate studies. It is not a stretch to say I would not be here today without their support, academically, financially and otherwise.

During my studies at SIAT, I found support and mentorship from Dr. Janet McCracken and Jim Bizzocchi. Also, fellow SIAT graduate students Brad Paras, Daniel Ha and Desmond Mulligan helped me to define exactly what being a graduate student is about, although I'm still not sure I know. From Dr. Goodman's lab, I would like to especially thank Nori Bradley, Ian Williamson and Jeremy Hamm. They provided me with much needed feedback and the occasional lesson in statistics. Also from Dr. Goodman's lab I would like to thank Jill Meanley, Eric Lee and Kirsten Willms for helping out with the *Heads Up Hockey* studies and indulging my occasional madcap requests. A final thank you to everyone I have met through Dr. Goodman's lab, you folks provided many memorable moments (Brent and Little D, I'm looking at you).

It was a pleasure to work with the *Heads Up Hockey* development team of Aaron Wylie, Glendon Holst, Craig Huff and Brad Paras. I am still amazed at the quality of work you were all able to produce without commercial production assistance. Obviously, without you folks this thesis would not be possible.

I am thankful to the fine folks at CIHR, SAGE and the SIAT committee that continually believed I could be an able teaching assistant. Without the financial support from these sources, not only would this thesis not be possible, I would not be possible. Some of my time and Dr. Mike Dobson's time was provided by the Social Sciences and Humanities Research Council (SSHRC) project: Simulation and Advanced Gaming Environments for Learning, which is a part of the Initiative on the New Economy and a Collaborative Research Initiative led by Dr. David Kaufman and Dr. Louise Sauvé.

Finally, to J.P., thanks for breakfast, lunch and dinner, as well as your amazing level of tolerance towards yours truly. The last two years weren't easy on me, but I bet they were more difficult for you, having to deal with me. Thanks for being there and thanks for being my sounding board.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF CODE SEQUENCES

GLOSSARY

Adventure game	This genre focuses on puzzle solving within a narrative framework relying on the player's ability to think logically.
Artificial intelligence[1] (game)	Refers to techniques used in computer and video games to produce the illusion of intelligence in the behaviour of non-player characters.
Behaviourism	An approach to psychology based on the proposition that behaviour can be researched scientifically without recourse to inner mental states.
Blog	A web-based publication consisting primarily of periodic articles (normally, but not always, in reverse chronological order).
Cognitivism	The approach to understanding the mind which argues that mental function can be understood as the 'internal' rule bound manipulation of symbols.
Concussion	An injury to the brain usually caused by a blow to the head, but not necessarily involving a loss of consciousness.
Constructivism	A set of assumptions about the nature of human learning that values developmentally appropriate, teacher-supported learning, often initiated and directed by the student.
Cutscene	A sequence in a video game over which the player has no control. Cutscenes are used to advance the plot, present character development, and provide background information, atmosphere, dialogue and clues
Digital game	A system in which players engage in an artificial conflict, defined by rules, that results in a quantifiable outcome.
Edutainment	A hybrid mix of education and entertainment that relies heavily on visual material, on narrative or game-like formats, and on more informal, less didactic styles of address.

[1] *c.f.* the research area that models human intelligence using logic languages

Gameplay Includes all player experiences during the interaction with game systems, especially formal games.

Karma The general definition is a measure of all that an individual has done and is currently doing.

LAN party A temporary, sometimes spontaneous gathering of people together with their computers, which they connect together in a local area network (LAN) primarily for the purpose of playing multiplayer computer games.

Simulation A computer program which simulates an abstract model of a particular system.

CHAPTER 1: INTRODUCTION

The following pages will discuss the design, implementation and assessment of an educational sports-action video game. The game, termed *Heads Up Hockey*, was specifically designed to modify the game playing behaviour of its participants by embedding an implicit teaching mechanism within the gameplay. The game was created as part of a multi-year CIHR grant[2] addressing concussion injuries in youth hockey. While a number of people were involved in this project, I had an instrumental role in the design and implementation of the game itself. In my role as game development lead my major responsibilities were to manage the development team, design and test the implicit teaching mechanism, and assist with the implementation where necessary. It is these experiences that form the essence of my thesis.

There were several motivations for selecting the video game as our educational medium over other media forms, such as video. The primary rationale was that we were dealing with youth hockey players who were of the age where video gaming occupied much of their time. The idea was to try and capture their attention through the video game medium that was specific to the sport in which they were participating and getting concussed. By embedding the teaching mechanism implicitly rather than explicitly, the goals of the game, at least to the participant, remained the same as in all games—to win. That is,

[2] Canadian Institutes of Health Research, Grant # CAR-42276, D. Goodman, P.I.

participants tried to outscore the computer-controlled opponent by skillfully playing the video game; there was no intent on behalf of the player to learn about concussions.

Using my unique perspective of leading the game development and the subsequent evaluation of the implicit teaching mechanism, I hope to present a coherent picture of our novel approach to educating youth regarding concussion in hockey. The foremost innovation of this game is an implicit teaching mechanism that was designed to lessen aggressive and negligent behaviours that could lead to concussive injury. The objective was for the player to be able to learn through playing the game, as opposed to learning explicitly from the game. Since these behaviour changes were designed to be achieved through non-obvious means, we were able to maintain a high level of engagement with the participant. We assumed that by keeping the player engaged with the gameplay for extended periods we would increase the time the learner is engaged with the material and the overall interest level with the activity.

Lessons learned from the implementation and assessment of *Heads Up Hockey* may be transferable to other educational situations. While our ultimate aim is to alter health related behaviours in the sports arena, the first step was to modify behaviour in the context of a sports-action video game. Further research will examine whether there is a transfer from video game playing behaviour to behaviour at the hockey rink.

In the chapters that follow, I first discuss some of the historical attempts at using games to teach and examine some of the common design features that are found in educational software. I find that a balanced mix of education and gameplay is not easy to achieve, despite the existence of several software titles that have been successful both as an educational tool and an engaging activity. In Chapter 3, several approaches for creating games within the education field are presented. These approaches hold promise to further the evolution of educational games away from a drill and practice model and into a role-playing experience. The history and design of *Heads Up Hockey* is discussed in Chapter 4. Important design constructs, such as how we operationalized our behaviour modification mechanism, are presented in detail. Chapter 5 describes a study that assesses *Heads Up Hockey* and its implicit teaching mechanism, and this is followed by analysis and results in Chapter 6. Finally, in Chapter 7, I present a summary and several conclusions based on my three year involvement with this project. Several possible improvements to the study and game will also be presented.

CHAPTER 2: DIFFICULTIES IN USING GAMES TO TEACH

This chapter explores several of the problems that face educational game designers. A background on the history of educational games and simulations is presented. Several themes emerge from this analysis that point to a common problem that many educational video games suffer from. That is, the problem of an indirect relationship between gameplay and content. To further illustrate this difficulty, several examples of educational video games that achieve an effective equilibrium between gameplay and content are presented, in addition to several others that do not. Finally, this issue of "Education vs. Gameplay" is explored in depth.

2.1 Background on educational games and simulations

Video games are receiving broad attention, both as an entertainment medium and as a research field. Digital games have survived attacks upon their value in society and have emerged stronger and more profitable than ever (ESA, 2005; Haddon, 1988). No longer do game researchers focus on validating their field with statistics about growing gamer populations and faltering movie industry revenues, rather they are now advancing the level of understanding about the video game phenomenon (Egenfeldt-Nielsen, 2005; Salen & Zimmerman, 2004; Squire *et al.*, 2003). Several new journals, books and conferences all testify to the acknowledged need to examine digital games from all sides in order to evaluate their impact on society and uncover their potential for more.

There are several areas that have received recent attention from games researchers, including identifying gaming genres, annotating game-development processes and the role of gender in gaming. Of these many areas, one of the more widely discussed topics is how to leverage video games for the purposes of education (de Castell & Jenson, 2003; Gee, 2003). It is not surprising that an academic field would look for ways to use the engaging power of digital games for education. Similar efforts have been made in the past to harness the success of entertainment mediums for educational purposes, such as film, television, audio recordings and so on (Saettler, 1968). With digital games, the challenge, like film and television, is to take the media-specific methods for creating entertainment and education and create a single experience that optimizes the contribution of both.

The history of educational video games stretches back to before the advent of the personal computer in 1981 and, like the entertainment-driven games market, has seen several success and failures (Egenfeldt-Nielsen, 2005). One of the major problems with educational software is that it is easier to retrofit educational content on to an existing game rather than design a novel and holistic educational experience. This results in a gameplay experience that can feel disjointed and ultimately nowhere near as engaging as a purely entertaining video game. The rest of this chapter explores some of the failures and successes of educational games and simulations.

One of the most powerful drivers of educational games and simulations is the world's armed forces. Since the 18th century, militaries have been experimenting with games to simulate battlefield strategies (Egenfeldt-Nielsen, 2005). This broad stroke use of games has narrowed over the years from high level strategy games to include specific skills training of individual soldiers and pilots. In particular, the US military is a leader when it comes to using video games for training. Whether training individual personnel or conducting theatre wide war simulations, the US military has found many ways to successfully employ video games (Prensky, 2001).

Military and educational institutions have the goal of transferring content and skills to a trainee in an efficient manner. Both institutions have approached games as a means to enhance their instructional programs. However, the military has claimed the most successes thus far in employing games to teach (Macedonia, 2002). The problem with some military games research is that it is subject to classification and may not be accessible by the academic community. The research that has been published has proven useful, at least in terms of advancing game technology (Prensky, 2001), but needs to be interpreted with caution. The military's educational gaming program has a large budget and an emphasis on relatively simplistic skill acquisition or combat training (Prensky, 2001). Non-military applications often do not have the same level of success due to lack of funds or having abstract educational goals, such as high level mathematics or programming.

Corporations and business schools have also advocated the educational use of video games. Simulations have been popular amongst corporations, from simple economic forecasts to more complex marketplace landscape predictions (Bertsche *et al.*, 1996; Dobson *et al.*, 2004). In addition, some companies use games for the purposes of managerial training (Prensky, 2001). However, these corporate training games often suffer from portraying a simplified real-life environment and high learning costs (Kyrylov *et al.*, 2004). Business schools have slowly increased the penetration of games into their curriculum and they are used today as often as case studies (Faria & Wellington, 2004). The moderate penetration of business-themed games into schools and corporations has not translated over to the commercial market. Despite the success of several business-themed games (e.g. *Rollercoaster Tycoon*) they are significantly out sold by shooting or action based games (ESA, 2005).

While business and military interests have had an influence on today's educational video games, nothing has guided the evolution of educational gaming more so than educational media (Egenfeldt-Nielsen, 2005). With the bulk of research that exists in other educational media, such as television and film, it is natural that those findings would initially be applied to video games. This has resulted in the production of educational video games that rely on principles learned from examining one-way media. Good educational results have been achieved using television and film, however, using these results to inform an interactive medium neglects the two-way (or multi-way) nature of digital games.

One result of this focus on the educational media perspective is 'edutainment'. Buckingham and Scanlon (2003) describe 'edutainment' as a "hybrid mix of education and entertainment that relies heavily on visual material, on narrative or game-like formats, and on more informal, less didactic styles of address." This heavy reliance on visuals and narrative results in edutainment games that have little connection between gameplay and what the learner is supposed to learn (Squire & Jenkins, 2004). Nonetheless, edutainment titles continue to dominate the educational gaming market (Egenfeldt-Nielsen, 2005), with little innovation seen from commercial vendors. On the other hand, the academic games community has produced several titles (e.g. *Virtual U*, *Making History* and *Revolution*) that attempt to expand the capabilities of educational video games.

The question is why large software companies are continuing to saturate a shrinking educational software market (Richtel, 2005) with software that contains little innovation. The roots of today's commercially available educational games date back to the 1960s. Educators were, and continue to be, eager to increase the interaction between learner and content, as demonstrated by the plentiful number of educational film and television productions. The 1960s saw a similar movement with the advent of computers, led by a massive investment in Computer Assisted Instruction (CAI) where the learner would perform the same task repeatedly until it was performed correctly. However, CAI lacked intrinsic motivation for the learner due in part to the lack of a "fantasy" element (Malone, 1981). To this day, the CAI legacy can be found in commercially available educational computer game titles, such as the *JumpStart* series and *Math Blaster* (Miller, 2000). However, with the pure 'entertainment' video game sector producing rich role-playing experiences, the educational game market can no longer compete. Ideas that are being developed within the academic community are needed to reinvigorate the educational game genre in order to fulfill its original promise of increasing learner interest, motivation, and retention (Hogle, 1996).

The number of journal articles, books and websites dedicated to digital game based learning has recently increased (Mitchell & Savill-Smith, 2004, pg. 1). The following is a selection of the research that Mitchell and Savill-Smith (2004) identify:

- descriptions of the use of computer games
- explorations of how different game attributes affect users' preferences
- overviews of the educational potential of video games
- articles based on pre- and post-testing of users playing educational games
- studies of the relationship between the use of games and lifestyle choices.

In addition to research, the academic community has also produced, or had a hand in producing, some of the more effective educational games and simulations (e.g. *LOGO*, *Dexter*, *Bronkie the Bronchiasaurus*, *Green Globs and Graphing Equations*, *Kristen's Cookies*, and *Virtual U*). Academic institutions continue to produce games, with organizations such as *Education Arcade* producing *Revolution*, *Environmental Detectives* and *Mystery at the Museum* (Education Arcade, 2005). While games produced by academic institutions often lack the funds needed for high-end production, they usually compensate through innovative instructional methods.

2.2 Examples of educational video games

One of the most popular and successful educational video games was first released in 1974. *Oregon Trail* by MECC remains a best selling educational title today. It was successful because the challenges in the gameplay were linked to what the player was supposed to learn. Everything from geography, to history, and even a little medicine was woven into the gameplay. Another example of a game that combined educational content with stimulating gameplay was *Trans Canada*. In *Trans Canada* the player assumed the role of a truck driver who had to deliver various goods to a variety of Canadian destinations. Again, Canadian geography, resource location and personal health management (sleeping and eating) were all integral parts of successfully playing the game.

Another classic example of a game that effectively integrated educational content and gameplay was *Rocky Boots*. The game is actually more akin to a simulation, and was described as "a visual simulation that made it possible for upper-grade-school students to design simple digital logical circuits, using a joystick to move around circuit symbols on the screen and plug them together" (Robinett & Grimm, 2005). This title sold upwards of 100,000 copies, partially due to its success of interweaving the educational content with the gameplay.

By 1977, software titles began to exhibit drill-and-practice traits that are the signature of edutainment software. *Basic Math*, followed by the similar *Electric Company Math* in 1979, showcases drill-and-practice principles with the addition of some simple gameplay. The 1980s saw the educational software market rapidly expand, competing with purely entertaining games for the consumer's dollar. The problem is that increased commercial interest often conflicts with developing effective educational games, since it is often easier and cheaper to produce a formulaic edutainment title. The commercial educational software industry realized they could make a profit without spending money on innovation and thus we see the start of edutainment's dominance through the late 1980s.

However, the educational software market saw its share of quality titles throughout the 1980s. Edutainment was coming, but it took some time to completely infiltrate the market. The adventure theme became prevalent in educational gaming, which produced games like *Snooper Troopers*, *Where in the World is Carmen San Diego?* and the venerable *Oregon Trail*. These games allowed the learner to learn through simply playing the game. The principle these games follow is to provide a meaningful play experience through which a player can learn.

One problem with the successful educational games in the 1980s was that they mostly relied on the adventure game genre. There is only so much adventure games can be made to teach, yet this theme continued through the early 1990s, producing titles such as *Winnie the Pooh in Hundred Acres Wood*, *Mickey's Space Adventure* and *Troll's Tale.* However, not all curricula can be taught through the adventure genre, thus other methods of instruction had to be explored.

In 1987, *Mavis Beacon* sought to teach players how to type by challenging them to complete certain letter combinations before time expired. It is a simple concept, rehashed in Sega's *The Typing of the Dead*, which pits the player against a horde of zombies. The learner must type the correct key combinations before the zombies get to them, otherwise their virtual brains will be eaten. This is a classic example of how edutainment separates gameplay from the educational content. Learning how to type and defeating hungry zombies have little to do with each other.

Through the 1990s, edutainment begins to push out other forms of educational software. Driven by profit margins, edutainment titles become more and more formulaic, relying on sequels to best sellers. Little innovation is seen, despite the best efforts of *Lucas Learning* and *The Learning Company*. *Lucas Learning* had the idea to create "a thoughtful game that's actually fun and helps kids to learn within the game medium" (Blossom & Michaud, 1999). This principle is everything that edutainment was not, yet had already been achieved in the past with the adventure themed games of the 1980s such as *Oregon Trail*.

A lack of innovation and growing consumer scepticism towards edutainment began to shake the once promising educational software market (Leyland, 1996). In response, game publishers tightened budgets, relied more on formulaic titles and forwent innovation. This vicious cycle of falling development funds leading to poor titles leading to further reductions in development funds contributed to the increasing sluggishness of the educational software market. Game publishers could make more money by investing in commercial game development, to the detriment of educational software. Eventually, educational computer games received a bad reputation and were no longer in high demand.

However, this had an interesting side effect. Certain games that were not intended to be educational ended up being so. It makes sense that with the number of games under development at least a handful would stumble upon a formula that mixed fun gameplay with some valuable educational content. The most famous examples from the early 1990s still hold enormous influence today: *SimCity*, *Lemmings*, *SimEarth* and *Civilization*. *SimCity* spawned many sequels on its own, and was the precursor to the latest "accidentally educational" software title, *The Sims*. The publisher of *Civilization* has released the fourth instalment of the epic turn-based strategy game, 14 years after the introduction of the original.

What these game designers created, and what makes educational software effective, is the player's ability to immerse themselves into a role. When a player dedicates themselves to a role, the challenges they encounter during gameplay can become learning opportunities (Leyland, 1996). For example, in *SimCity*, if a player does not understand the side effects of a coal power plant, they would eventually learn first hand that coal is a polluter. Similarly, with *Civilization*, if a player does not understand the difference between a Monarchy and a Democracy, they can not effectively rule their world. Winning these games requires learning the causal relationships between concepts, such as land value and taxes (*SimCity*) or government and military (*Civilization*).

Through the mid 1990s, there were efforts to create educational games that effectively mixed gameplay with educational content, but edutainment titles continued to dominate. However, it soon became apparent that the edutainment formula was only effective, commercially and otherwise, for elementary and pre-school children. The realization of this smaller demographic has reduced the optimism previously associated with games and learning (Richtel, 2005).

Even today, some of the best educational titles on the market have a pedigree that date back to the mid 1980s. *Oregon Trail* is still on the market, as well as a host of *Electric Company Math* clones (e.g. *Math Blaster* and *JumpStart*). One current trend is to associate a brand name, such as *Disney*, to the edutainment software, thereby giving it some credibility and hopefully some marketability. However, there remains a large untapped potential for the development of quality educational games.

This untapped potential is being explored through the production of video games by academic institutions and companies dedicated to learning software. For example, *Muzzy Lane* is a software developer and publisher dedicated to creating multiplayer educational gaming solutions that complement and improve on traditional teaching tools (Muzzy Lane, 2005). They created the simulation game *Making History* that covers the causes and consequences of World War II. The game puts the player in the role of a head of state, leading a struggling nation through challenges based on real historic events, using accurate historical data. For each scenario there is a description of the main events and challenges, a list of key themes, and a list of post-scenario questions that an educator might use to help students analyze and learn from the results of playing that scenario. The key to *Muzzy Lane*'s approach in *Making History* is the situated aspect of the gameplay. The player relives the history through role play, rather than through a separate history exercise that is presented after some unrelated gameplay.

Academic institutions are also key innovators in educational gaming technology. An example of such an institution is the Academic Advanced Distributed (ADL) Learning Co-Lab based out of the University of Wisconsin-Madison. Their goal is to serve as the focal point for academia in promoting high quality, reusable content for distributed learning. They produced a prototype simulation that allows users to experience and apply the first five steps of the CDC established "Outbreak Detection Process," immersing the player within the simulation's narrative (Martinez-Gallagher & Norton, 2004). Similar to other successful educational games, this simulation situates the learner within an environment that teaches its content as the participant explores the world.

Academic institutions, and in particular medical schools, have produced several games designed to improve self-care health behaviours (Lieberman, 2001). An example from this genre is *Packy & Marlon*, produced by the now defunct *Click Health*. *Packy & Marlon* is designed to improve a young person's ability and motivation to undertake the rigorous self-care necessary to control insulin-dependent diabetes (S. J. Brown *et al.*, 1997). The player assumes the role of Packy, an elephant with diabetes. During the game, the player must monitor their character's blood glucose if they are to save their camp from rats that have scattered the food and diabetes supplies. Brown *et al.* ran a controlled trial to test the effectiveness of the game. The results showed that the game significantly improved diabetes-related self-efficacy, communication with parents about diabetes and self-care behaviours while significantly decreasing unscheduled urgent doctor visits.

In addition to educational gaming, academic researchers have long been active in the simulations[3] community, lauding the benefits of simulations for more than 40 years (Wolfe & Crookall, 1998). One example of an innovative simulation game is KM QUEST (Leemkuil *et al.*, 2003). This simulation game is designed to teach how to solve knowledge management (KM) problems. Leemkuil *et al.* claim that the basic instructional strategy used in simulations and games is constructivist. However, despite the recent shift from instructivist approaches toward constructivist approaches (van Merriënboer, 1997), Leemkuil *et al.* do not discount the role instructivist strategies should play in instructional design. The optimal learning environment blends instructivist and constructivist approaches based on the learning goals, types of problems students have to solve, prior knowledge of the students and the context in which learning takes place.

---

[3] The difference between a simulation and a game is subtle. Jacobs and Dempsey (1993) state that the distinction between the two is often blurred, even though they are not the same (Sauvé *et al.*, 2005). Many articles now refer to a single simulation game entity.

2.3 Education vs. Gameplay

The most successful educational games find an equilibrium between gameplay and content such that the learning tasks are perceived by the player as a true element of the gameplay (Fabricatore, 2000). However, achieving such a fusion is a complex proposition and is one reason why so few good educational titles exist. Recently, groups have been formed (e.g. *Education Arcade* and *Serious Games*) to explore what kind of balance is needed to create effective educational games. One of the strategies that these groups promote is to combine commercial game developers, educators and subject matter experts in the creation of educational games. The idea is sound, but there remain significant problems to overcome. The challenge is to make the disparate goals of the various groups converge as much as possible.

Looking back on the experience of implementing *Heads Up Hockey*, part of the problem was that the educators did not understand the "game" language used by the development team. There was a lack of a common vocabulary that could explain why a certain idea might harm the playability of the game. Similarly, the educators often could not put into understandable terms for the game developers why their ideas were important to include. In the end, our project focussed on gameplay more than traditional educational principles because we wished to employ more contemporary learning designs that are more tailored for video game use, such as role-playing. Ashley Lipson in Prensky (2001, p. 152) would agree with our approach, saying, "To be an entertaining and educational game, it must first be a game, and only then, a teacher."

Creating a game that optimizes the mix of gameplay and instruction is a complicated task, due in part to the lack of regular communication between the various disciplines needed to create an educational game. Unfortunately, communication is not the only hurdle in creating educational software. Even if the team works well together, they still must come up with a design that will both engage and educate. Another problem is that most game developers become game developers because of their love of video games, not education. Despite reports of how rewarding educational game design can be (Prensky, 2001), most top end game developers prefer to work on purely entertaining digital games. This "brain vacuum" is a serious hindrance to the creation of quality educational titles. Finally, educational titles receive relatively little money for innovation and cutting edge technology, often resulting in the production of narrowly used research tools or formulaic edutainment (Squire, 2004, pg. 66).

In summary, the "Education vs. Gameplay" issue is a problem of "indirection". Indirection within an educational game is the result of a design that disassociates the gameplay from the content, which Rieber (1996) calls an "exogenous fantasy". For example, a fictional educational game could reward a successful match of Tic-Tac-Toe with a lesson on Canadian History. Playing Tic-Tac-Toe is meant to extrinsically motivate learners to absorb the Canadian History content. However, the relationship between Canadian History and Tic-Tac-Toe is bound by a behaviourist mechanism that artificially relates the game to the content. There are several real world examples of indirect educational games, including *Math Blaster* and *The Typing of the Dead*, but few scientific studies on their effectiveness. These games and many others suffer from the problem of indirection, with the result being neither entertaining nor educational. The solution to indirect educational games will be explored in the coming chapters.

Although not numerous, there are several studies that support the premise that indirect educational games are ineffective at teaching. McMullen (1987) investigated the effect of informational, drill, and game format computer-assisted instruction (CAI) on the achievement, retention, and attitude toward instruction of sixth-grade science students. An informational CAI lesson on Halley's Comet was administered to three randomly selected groups of sixth-grade students. A CAI drill about the content of the informational lesson was given to one group, and a CAI game was given to another group; only the informational lesson was presented to the third group. No significant differences were found between the groups on a post-test measuring achievement given immediately after the instruction or on a retention post-test given one month later.

Din and Caleo (2000) investigated whether kindergarten students who played the *Lightspan* console learning games learned more than peers who did not play such games. The *Lightspan* learning games suffer from indirection in that the goals of the game (e.g. to teach math) are not directly related to the gameplay. The experimental group played the games for 40 minutes per day in school for 11 weeks. Findings from the data analysis indicated that the experimental group performed no differently than the control group in the math area. Although this study did produce some positive results in areas that have historically responded well to drill and practice models (e.g. spelling), the authors concede further research must be performed to say that playing the *Lightspan* series leads to learning.

CHAPTER 3: APPROACHES TO CREATING GAMES FOR LEARNING

The previous chapter charted some of the history of educational games. That history has been marked by a number of successes, but educational gaming is in danger of becoming an irrelevant software segment. New approaches are needed to reinvigorate innovation within educational game development. This chapter looks at several of the issues and approaches that are influencing the educational gaming field. The remaining chapters follow these approaches by showcasing one possible design and implementation of an implicit learning concept. There exists other game-based educational trends, such as the increased interest in games and health (Serious Games, 2005; Silverman *et al.*, 2002), the potential of massively multiplayer online games (Steinkuehler, 2004), the proliferation of academic research centres dedicated to gaming (Ludology, 2005) and the ongoing debate on the role of video games within the K-12 education system (Maushak *et al.*, 2001). However, I will focus on approaches that directly relate to our project and in particular those approaches that can help solve the problem of integrating gameplay with the educational experience.

3.1 Issues in educational games research

Using games and simulations to teach is not a new concept and as such there is a large body of research on the topic. There are several themes within the games research community that relate to the design of *Heads Up Hockey*, including: a) play and pleasure, b) game-development, systems, and content points of view, c) narrative and gaming, d) psychological, behavioural, and cognitive effects of gaming and e) constructionist theory and research (de Castell & Jenson, 2003). The goal of much of this research is to determine how to take the best elements of commercial video games and meld them with an effective instructional program. In practice, this is often a very difficult task given the somewhat divergent goals of entertainment and education. To illustrate some of these difficulties, de Castell and Jenson outline some of the development challenges facing the *Ludas Vitae* project, which "seeks to bring play and education together in an ecology-focused, character-driven, online garning environment" (de Castell & Jenson, 2003, pg 658). A sampling of their challenges include: a) employing 3D rather than 2D graphics to satisfy the expectations of young learners, b) creating a set of rich, fully navigable immersive environments, c) designing gameplay to fit into a first-person perspective, d) ensuring equal rewards for skill development as for content mastery, e) embedding all learning within the narrative of the game, and f) allowing the learner to assume diverse characters.

While these challenges will not apply to all development efforts, the *Ludas Vitae* project and *Heads Up Hockey* share several common goals. The issue of 2D versus 3D effects not only the learner's graphical expectations, but also the level of simulation fidelity and its effect on learning outcomes (Feinstein & Cannon, 2002). It is not necessarily the case that a higher fidelity simulation will result in greater learning performance, due to the possibility of the learning goals being obscured by tangential information (Thiagarajan, 1998). It is important to understand what applications best suit 3D or 2D graphics and what each paradigm affords the instructional designer (Cockburn & McKenzie, 2002). In addition, *Heads Up Hockey*, like the *Ludas Vitae* project, pays special attention to the divide between skill mastery and content mastery. The common challenge is to intertwine the process of skill mastery with content acquisition. How to actually achieve this intertwinement between gameplay and content within a simulation or game world remains a difficult problem.

Creating an effective intertwinement of gameplay and content is one component of the broader challenge of creating an instructional environment. Reigeluth and Schwartz (1989) have addressed the problem of designing instructional simulations and they identify five key features of simulations that act as vehicles for achieving acquisition, application and assessment: 1) generality, 2) example, 3) practice, 4) feedback and 5) help. "Generality" refers to system generated statements that provide insight into the higher order organization of the procedures and principles being explored by the learner. As later sections will describe, all of these features are reflected within the design of *Heads Up Hockey*, although some features were more well-defined than others (e.g. practice and feedback more so than help and generality).

Thiagarajan (1998) is another researcher that has addressed the problem of creating and deploying instructional simulations. Several of the guidelines that Thiagarajan lists can be applied to the design of *Heads Up Hockey*, such as:

- Be wary of what people may unconsciously pick up from participating in your simulations
- Carefully select the level at which you want to simulate a system
- Don't constrain yourself to one type of simulation
- Don't limit the range of application of simulations
- Select the appropriate level of fidelity to suit your needs
- Use a variety of design approaches for creating simulations

These guidelines are not a prescriptive methodology for creating an effective simulation and are vague in some cases, but they direct the designer to ask important questions that many omit.

Another facet of games research focuses on the power of games to engage and motivate players to overcome obstacles that arise during the course of gameplay. Malone (1981) described the elements of a "fun" game and defined a rudimentary theory of intrinsically motivating instruction that is based on 1) challenge, 2) fantasy and 3) curiosity. Garris, Ahlers and Driskell's (2002) review of the literature expands upon Malone's characteristics of an intrinsically motivating learning environment by adding 4) rules/goals, 5) control and 6) sensory stimuli. Malone's and Garris et al.'s goal was to discover the essence of what makes games and simulations engaging and then re-apply that model to an instructional setting. However, some of these characteristics are easier to design and implement than others. Poorly designed edutainment may still be challenging, allow user control, present the learner with a set of goals and evoke sensory stimuli, but they often do not provide an engaging fantasy. The importance of fantasy is furthered by Rieber (1996) through his classification of fantasy into two types: endogenous and exogenous. An educational game that employs an "endogenous fantasy" weaves the content into the gameplay, whereas "exogenous fantasies" can be thought of as educational "sugar coating". Creating an effective endogenous is emotionally appealing, can motivate and engage, and can lead to greater learning (Asgari & Kaufman, 2004).

It is clear that there are many opinions on what makes an effective simulation or game, but how do we measure this effectiveness? Feinstein and Cannon (2002) developed a framework for pursuing the evaluation problem by considering three major constructs: fidelity, verification and validation. Fidelity and verification address the level of realism and the degree to which the model performs as intended, respectively. However, the most important aspect of simulation evaluation is validity, that is, the correctness of the model. Feinstein and Cannon identify two types of internal validity important to the evaluation of simulations: representational validity and educational validity. Representational validity asks to what extent a simulation game accurately represents desired phenomena. Educational validity asks to what extent are student decisions influenced in the intended manner by game design. For the purposes of *Heads Up Hockey*, both educational and representational validity are important concerns. The accuracy with which we capture the game of hockey (one aspect of representational validity) is central to a player's immersion within the game, and hence their intrinsic motivation to learn. Influencing student decisions through the game design (educational validity) is the core purpose of *Heads Up Hockey*, thus its level of validity is a chief concern.

3.2 Approach: Bring people together

In addition to the issues mentioned above, one of the largest problems in creating an educational game is getting educators, subject matter experts and game developers to understand each other's priorities and perspectives. That is why there has been a surge in the founding of initiatives and organizations whose mandate is to bring these people together in the name of improving educational games. *Education Arcade* and *Serious Games* explicitly list in their charters a need to bring interested disciplines together (Education Arcade, 2005; Serious Games, 2005). The hope is that by organizing conferences, hosting online discussions and encouraging research, disparate groups can begin to have a dialogue free from semantic misunderstandings. A significant amount of effort has gone into trying to specify and disambiguate the roles of educational software contributors, with the end goal being improved communication (van der Mast, 1995). Thus, it was promising to hear at the 2005 DiGRA Conference plenary in Vancouver that the communications hump was beginning to be overcome. While games researchers remain couched in vastly different backgrounds, each plenary speaker expressed how they were encouraged by the apparent increase in common understanding of the history and problems of games research.

There is hope that a partnership between Microsoft and MIT will lead to more innovation in educational games. Dubbed the *Games-To-Teach* project, it had a charter to "develop conceptual prototypes for the next generation of interactive educational entertainment" (Games-to-Teach, 2005). It has now been amalgamated with the broader *Education Arcade* project. One of the criticisms levelled in the previous chapter was lack of funding and access to the best game design minds. This partnership of academia and industry, under the *Education Arcade* banner, helps to partially alleviate the funding issue. However, since Microsoft and MIT are not game developers by trade, they may not have easy access to high end game developing talent. Nonetheless, the spirit of bringing people together is an approach that gives hope to the future of educational software.

Finally, federal grants to explore the usefulness of games in education are becoming more common (NSF, 2005; SSHRC, 2005). The grants often stipulate that the work be done over a wide range of groups, challenging them to come to some sort of collective working platform. Governmental grants are an important step to move educational gaming into the next generation, as most commercial game developers are locked in an ultra competitive games market. Development firms have little time and money to fund innovations in educational gaming, therefore it is now the government's prerogative to step in and fill this funding void.

3.3 Approach: Incorporate the social context

Another approach that could improve the effectiveness of educational games comes from Egenfeldt-Nielson's (2005) generational perspective. According to Egenfeldt-Nielson, educational games can be split into three generations, with the third generation still not having completely arrived. The first generation is edutainment, which is based on behaviourism and has been roundly chastised for being ineffective save for a small niche of applications. The second generation focuses on the learner and is influenced by cognitivism. Here, there is more space for the learner to explore solutions on their own terms, but the experience is still situated between the computer and the player. The third generation expands on the second by taking into account the wider social context of game playing. Settings, facilitators and communities are all important for the future success of educational gaming (Egenfeldt-Nielsen, 2005; Kirriemuir, 2003).

One of the problems with educational games has been that simplified knowledge or behaviours (e.g. typing, basic arithmetic, and eye-hand tasks) have been most successfully transferred via video games. The thought was that as long as the visuals were engaging enough then the student would absorb the ancillary educational content. In their review of games and learning literature, Kirriemuir and McFarlane write:

> Rather than aiming for an experience that superficially resembles leisure-based 'fun' activities, or one which attempts to conceal the educational purpose, it might be argued that we should understand the deep structures of the games play experience that contribute to 'flow' and build these into environments designed to support learning. (Kirriemuir & McFarlane, 2003)

The key idea from this quote is to move away from gameplay that is tangentially associated with the educational purpose and start building learning environments that contribute to 'flow' (Csikszentmihalyi, 1992), or engagement. To accomplish this, the social context of the learning environment should not be ignored. Humans naturally learn in groups, and without communities, complex ideas and knowledge often cannot be learned (Lave & Wenger, 1991). Knowledge is not constructed just between player and machine, but between the interplay of the subject, the tool (game) and the surrounding community (Engeström, 2001; Squire, 2002).

One ramification of purposely incorporating the social context into educational games is the need to support communication outside of game playing time. For example, the game could be packaged with a handbook that includes suggestions on how to connect game experiences to concrete real world examples, how to encourage a community of practice and how to increase communication amongst learners and teachers. Another important aspect of incorporating the social context into educational games is a well-versed facilitator (Kirriemuir, 2003). The facilitator needs to have intimate knowledge of the game as well as the subject matter. If the facilitator understands the problem that the learner is experiencing they can quickly provide a solution and provide a more seamless learning experience. This combination of video games and facilitators has been successfully attempted before, often using commercial software such as *Civilization III* (Squire, 2004) or *SimCity 2000* (Adams, 1998).

The 3rd generation of educational games is more suited for structured environments like schools, but there is also an application for the home market. Even though the social context is more difficult to integrate when a child learner plays on their own at home, there are still ways to develop the social context. First and foremost, if the parent can take on the role of the facilitator this could greatly increase the effectiveness of the game. However, parents have typically shied away from getting heavily involved in their child's game playing, often because they do not play games themselves. Thus, if the home experience can not support a facilitator role, the game should encourage and incorporate online multiplayer elements. While not as effective as a face to face experience, the game can still emphasize the social context by using the Internet as the communications conduit between like-minded learners. These online communities could have dedicated experts (facilitators) that would serve the same purpose as the teacher in a school. Finally, the game could be designed such that a player is encouraged to invite their friends over to participate. History tells us that if a game is engaging enough, the players will naturally form their own gaming "parties", thereby providing a face-to-face social context (Jansz & Martens, 2005).

3.4 Approach: Implicit learning

Implicit learning is "the process by which knowledge about the rule-governed complexities of the stimulus environment is acquired independently of conscious attempts to do so" (Reber, 1989). Designing educational games to utilize our ability to implicitly learn holds promise for easing the integration of gameplay and educational content. Implicit learning, as a game design principle, is similar to the concept of "stealth learning". A good definition of stealth learning is found in de Castell and Jenson (2003).

> '[S]tealth learning' [is where] players learn subliminally or incidentally through rule structures, tasks, and activities within the game. (de Castell & Jenson, 2003, pg. 655)

However, there are strong arguments against the idea of stealth learning that deserve acknowledgement. Okan (2003) and Kirriemuir and McFarlane (2003) identify researchers who are not eager to concede that learning is not fun and should be hidden from view. Another argument against stealth learning is that if the instruction is not obvious enough then there is little chance for reflection since the learner is unaware that they just learned something. This is a problem because reflection has been argued to be an important aspect of the learning process (Moon, 1999). In addition, since educational games rarely employ a multi-player model, the opportunity for group reflection is lessened. Therefore it is important for single-player educational software to allow time for the learner to reflect upon their experiences (Svane *et al.*, 2001).

There are two adjustments that can be made to stealth learning that may answer some of the criticisms. First, as previously discussed, Egenfeldt-Nielsen's 3rd generation of educational games stipulates that the game should not exist in a vacuum between the player and the computer. By expanding the social circle and including classmates, friends and facilitators, the possibility of reflection is introduced. Even though the learning may have been by stealth, by exploring each other's experiences in a community of practice, the learner will be able to reflect on a wide cross section of experiences. In this way, the salience of the learning only has to be as explicit as the most insightful member of the community. However, there is always the possibility that the most insightful member will be too shy to speak their mind in the face of a more vehement participant. Therefore, it is important that the most insightful member's views are encouraged and brought to the forefront of the conversation by the facilitator.

The other adjustment is to not take stealth learning to strictly mean that all traces of instruction must be hidden from the player's view, but as a principle that guides educational game developers to integrate gameplay with content. In other words, the stealth learning game designer should concern themselves with designing a teaching construct that gracefully compliments gameplay rather than hiding all visible cues of education. The important question is whether participants learn, and not whether it was by stealth or otherwise. This means that if a particular context demands that the game present its content more obviously, then designers should be free to ease the stealth aspect of the instructional design. The difficulty is ensuring that the playability of the game is not affected by these design changes.

This relaxed perspective of stealth learning gives way to an implicit learning approach. The primary goal of an implicit teaching mechanism is to avoid the direct presentation of material that is incongruent with gameplay. Therefore, even if the participant does not have any intention to learn, through the primary act of playing an implicit learning game they will be able to acquire knowledge or behaviours. Implicit learning has a broad base of research to draw on within the cognitive science community, which can help inform the design of an implicit teaching mechanism. However, there is no widely accepted definition of the implicit learning process and researchers continue to work on this problem (Rose *et al.*, 2005; Shanks *et al.*, 2005).

Even though the details of the implicit learning process are largely unknown, there is little debate that the process is active in all healthy learners. For example, Bechara *et al.* (1997) designed a gambling experiment where participants were faced with four decks of cards. Two decks of cards carried an immediate reward of $100, with the other deck carrying rewards of $50. The decks that had the large $100 reward also carried a disproportionate number of penalty cards, which led to the $100 decks performing more poorly than the $50 decks. The results showed that participants began to choose advantageously before consciously grasping the correct strategy. The authors argue that overt reasoning is preceded by a non-conscious biasing step. If this is the case, then a designer could create a mechanism that engages this non-conscious biasing step in order to affect behaviour and create an opportunity for learning.

An implicit teaching mechanism solves the problem of 'indirection' that previous educational software efforts have suffered from by manipulating the context of play as opposed to the content of play. This approach creates an "endogenous fantasy", where the content is weaved into the game (Rieber, 1996). The "exogenous fantasy" in the Canadian History/Tic-Tac-Toe example (Section 2.3) employed an instructional design that attempted to change the definition of Tic-Tac-Toe to yield lessons on Canadian history upon victory. This is an example of changing the content of play in an attempt to teach. An example of changing the context of play would be to move the traditional learning context of Canadian History from didactic instruction to an immersive interactive environment. Inside this environment, we can further manipulate the learner's context to assume different historical roles or present alternative histories based on the learner's decisions. Through the learner's role playing they will be intrinsically motivated to absorb as much content as possible so that they can make informed decisions on how to further their standing in the game.

3.5 Research question

The question that this thesis attempts to answer is how educational digital games can be developed in a way that preserves inherently enjoyable gameplay aspects at the same time as reliably changing learner outcomes. Existing educational efforts often suffer from the problem of indirection, in that the gameplay and content are indirectly related to each other. This can result in an experience that is neither educational nor fun. To combat indirection it is important to ensure that a player can achieve the learning goals as a direct result of playing the game. In other words, the content should be implicit to the playing of the game. Thus, we based our instructional method on the concept of implicit learning in order to increase safe-play behaviours within the context of an educational hockey video game. The implementation and evaluation of this game is presented in the following chapters.

CHAPTER 4: THE DESIGN OF HEADS UP HOCKEY

This chapter discusses the background and history of our educational video game *Heads Up Hockey*. This chapter also details important design elements of the gameplay and educational content, and how they were integrated. Descriptions and screenshots of gameplay will be presented to illustrate the experience of playing *Heads Up Hockey*.

4.1 Background

Over the last several years, there has been a steady flow of research focusing on the dangers of mild traumatic brain injury (concussion) in sport, especially amongst young athletes (McCrory *et al.*, 2004). Our particular concern is with educating youth hockey players about concussion. Concussions have often been underreported, which leads to youth hockey's administrative bodies having insufficient information for making decisions on policies designed to make the game safer (Williamson & Goodman, In press). In addition to education, there are other methods of concussion prevention that are currently being considered and employed (McIntosh & McCrory, 2005), such as protection via equipment. However, research indicates that concussion prevention is a problem whose solution does not appear to be achievable through protective equipment improvements alone (Curnow, 2003; McCrory, 2001). Certainly protection is one component in reducing brain injury, but education must play an important role in making hockey and sport in general safer.

Canada's national administrative body for hockey (Hockey Canada) has seen it fit to explore these issues. National granting agencies in both the U.S. and Canada have supported research into the effects, causes and prevention of concussion (Cook *et al.*, 2003). Current educational efforts often combine concussion with other injury prevention issues, such as the danger of spinal cord injury. A review of hockey related injury prevention educational material (Cook *et al.*, 2003) identifies four concerted efforts to educate hockey players about mild traumatic brain injury. The four efforts Cook *et al.* identified were *Smart Hockey with Mike Bossy* (ThinkFirst Foundation of Canada, 1988), *Heads Up: Head Injury and Trauma In Sport* (Sport Medicine Council of Alberta, 1999), *Heads Up, Don't Duck* (Massachusetts Medical Society Department of Public Health and Education, 2001), and *Heads Up Hockey* (USA Hockey). None of these programs have undergone systematic evaluation.

The educational program Cook *et al.* evaluated was ThinkFirst Canada's instructional video, *Smart Hockey.* The video demonstrated both safe and dangerous tactics by using real game footage and well known hockey personalities. They found an improvement in concussion knowledge as measured by a questionnaire administered before, immediately after, and three months after exposure to the video. There was no significant decrease in total penalty minutes in the experimental group, but body checking related penalties were significantly reduced. These are good initial results, but one significant problem with this video presentation is that it runs long, clocking in at one hour. Despite ThinkFirst Canada's success in producing an engaging film (Mats Sundin of the Toronto Maple Leafs is featured), it is an extended amount of time for a young player to focus their attention on an injury prevention video.

Educational media, such as video, can be effective but it is fast becoming a tool of the past. With the penetration of computers, television, cell phones and video games into homes and schools, younger generations now demand high levels of parallel interaction. This group of people has been dubbed the "Twitch Generation" (Lancaster, 1998), stemming from their ability to quickly process information from a variety of concurrent sources. Of course, the Twitch Generation encompasses today's young hockey players. Therefore, we hoped to leverage the intrinsically motivating aspects of interactive media to educate about the dangers of concussion in hockey. The approach we adopted was to create an educational hockey video game that attempts to alter the unsafe game play behaviours that can lead to concussions.

4.2 The history of *Heads Up Hockey*

At the start of this project we knew we wanted an interactive video game but the exact dimensions of the project were yet to be determined. The primary design parameter that we had was that it would be a 3-on-3 interactive hockey computer game. The first documented account of the concept of developing an educational hockey video game was in an initial CIHR letter of intent. The original letter of intent for the research project "Mild Head Injury in Youth Sport: There Really is Cause for Concern" was submitted in March 2000. The full proposal, further detailing the plans for educating youth about concussions via video games was submitted in September 2000, and fully funded in March 2001 for a five year term. The design and development of *Heads Up Hockey* began in January 2003 at which time I became fully involved.

The first several months had a lasting impact on the way the project would ultimately unfold. It was decided in late January 2003 to create the game in Macromedia Director. This decision would have lasting effects on the game's development because of the ease in which Director allows visuals to be manipulated, but the difficulty in which it handles programming complex structures. Once we had decided on the Director platform, the next major issue to resolve was whether to use 3D graphics or 2D graphics. 3D graphics are the *de facto* standard for modern games, but they also incur a significant amount of overhead in coding expertise and performance. On the other hand, even though 2D gaming graphics are no longer widely used, 2D animation sequences can be rendered from 3D models that maintain the 3D appearance. The other reason for selecting 2D graphics was a clearer and quicker implementation path that was more within our team's skill set. Thus, by March 2003, we decided the game would be done in 2D using Director.

By late March, we had a team that included a Macromedia Director programmer and a visual artist. The artist and programmer began to work through artistic style and user control issues while I began to design the physics for our simulated hockey environment. Each object interaction required a complex representation, from shooting, to passing, to hitting, to scoring. All of these events required some representation of a solid object interacting with forces and other solid objects and required several months of effort.

Figure 4.1: Early screen shot of *Heads Up Hockey* from May 2003

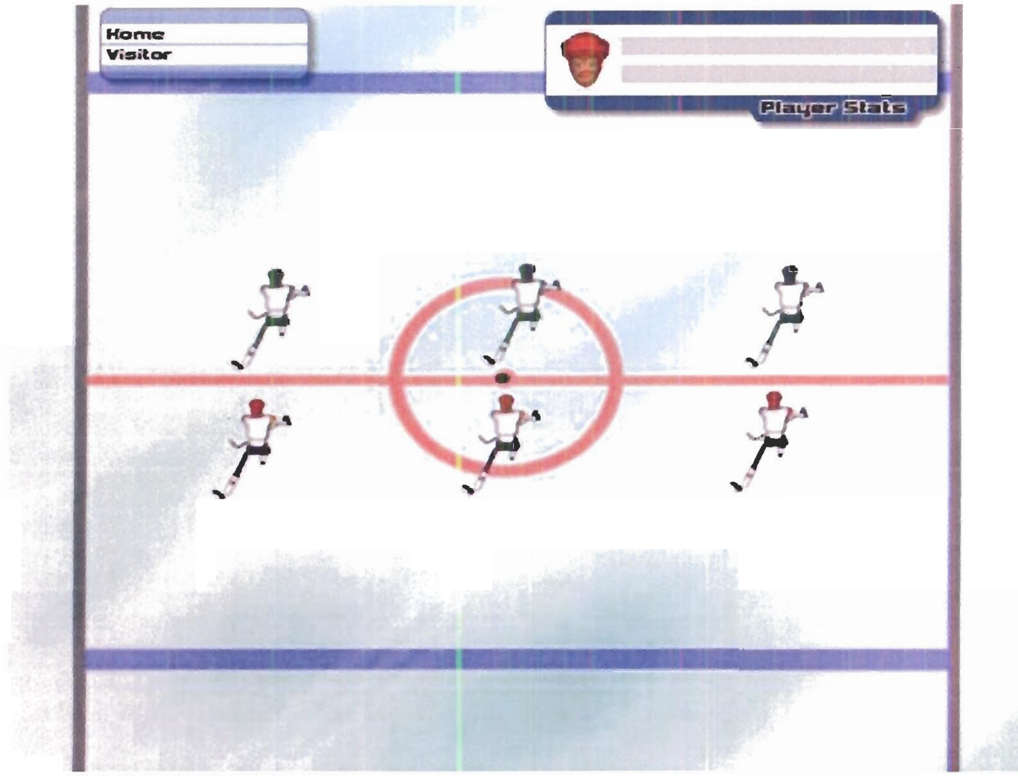

From *Heads Up Hockey* version 056 (current version is 332)

By mid-April 2003, we had finally settled on a name for the game, eschewing *Pressure Zone*, *Rockin' Ocky*, *Puck Style* and *Cross Check* for the eventual choice *Heads Up Hockey*. Around this same time, *Heads Up Hockey* also had a running demo. It consisted of no more than a few icons "skating" around on an ice surface, but you could definitely see the roots of what the game was eventually to become. There were no collisions or any other kind of collision detection, but there was a rudimentary shooting algorithm, which would eventually be rewritten.

Satisfactory progress on all aspects of the game continued throughout spring 2003. However, it soon became apparent that our team was missing the ability to create intelligent behaviours in opponents and team mates. To address this issue, we added a dedicated AI programmer/designer by the end of May 2003. This brought our development team size to 5 members, which is relatively small compared to more typical commercial development teams of 15 or more members.

The first major instructional contribution was the inclusion of animated cutscenes. Cutscenes are sequences in a video game over which the player has no control. They often are used to advance the game's plot, present character development, and provide background information, atmosphere, dialogue and clues. The key to this content delivery method was to create 'situations' within the cutscenes that were believable within the narrative of the game world. These situations were then used as opportunities to present the content. Cutscenes were triggered after a player received a simulated concussion, at which time concussion-related content was presented in the context of a conversation between players or a player and coach. The cutscenes were a believable extension of the gameplay, which holds players within the 'magic circle' (Salen & Zimmerman, 2004). To further the player's suspension of disbelief (Murray, 1997), gameplay disruption was minimized by only showing cutscenes during stoppages in play (e.g., after the puck has been shot into the crowd).

Figure 4.2: Screen shot of an animated cutscene

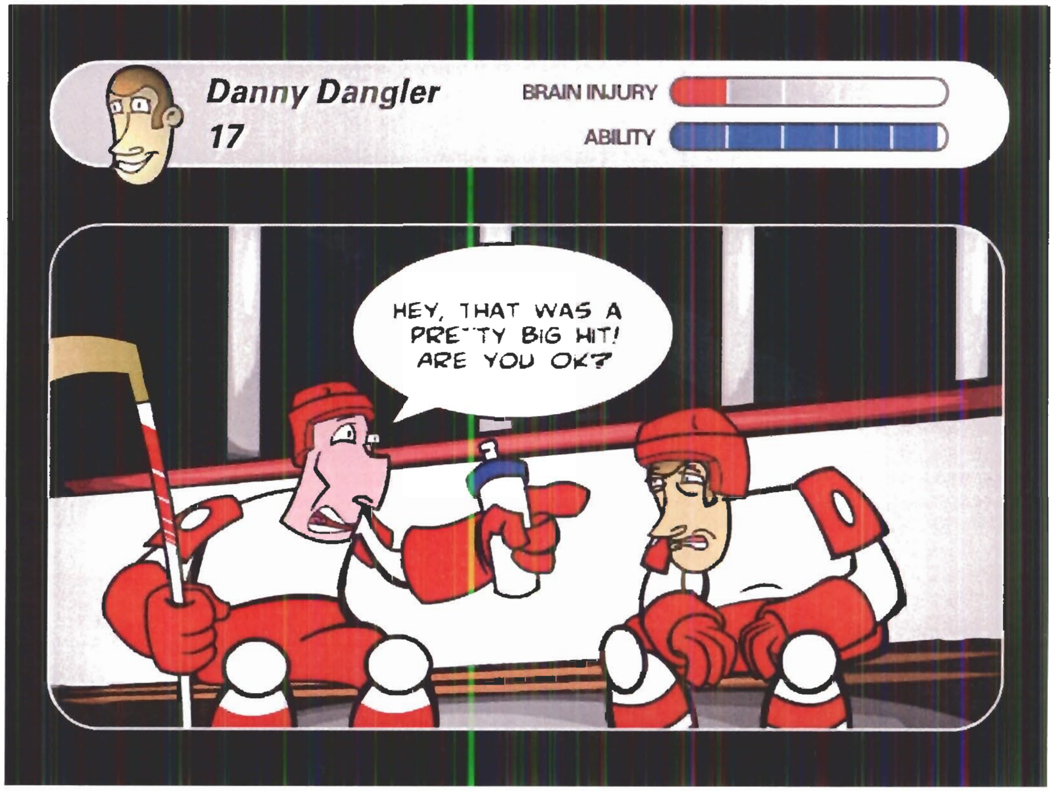

Cutscenes were not included in the version of *Heads Up Hockey* tested for this thesis

The specific teaching mechanism that is examined in this thesis was formulated in March 2003. The idea was to create a system that would penalize aggressive behaviour, and reward positive behaviour. In addition, the occurrence of a reward or penalty would not be directly revealed to the participant, thus allowing the learner to discover the connection between positive behaviour and reward on their own terms (Jonassen, 1997). Thus, the notion of 'karma', a system that would keep track of aggressive and negligent player events and subsequently dampen the player's ability to win, was born. Similarly, if a player exhibited positive behaviours, the karma system would reward the player with a greater ability to win. I believed karma would be an effective paradigm to change behaviours because it has roots in human nature. If we perform certain actions and those result in unfavourable outcomes, we will consider changing our behaviour (Burton *et al.*, 2004). This extends to the situations where actions and outcomes do not seem to be directly connected (Skinner, 1948). Superstition is one manifestation of our natural tendency to change our routines based purely on coincidental correlation between events, also known as the *post hoc* fallacy. For example, some hockey goaltenders touch their posts before each game because they believe that act will bring them good luck. By creating a karmic mechanism that is based on situated and constructivist learning principles, and takes advantage of our innate ability to correlate actions with outcomes, I hoped to induce behaviour change.

Over the summer of 2003, the project achieved several major development milestones, including the integration of a player collision system and the finalization of the player animations. There was an August deadline for testing the game with youth hockey players enrolled in hockey summer camps. Even though the physics, AI, graphics and overall playability had reached a technical level of integration, several playability and usability issues had to be addressed. Thus, the focus changed to running a pilot test that would give us feedback on the playability and flow of the game. This pilot test was not measured in any formal way, but we were able to observe and interact with participants in order to elicit their feedback on the game. The main outcome from that pilot test was that we had to make the game easier to win for novice users.

Near the end of 2003, we continued to work on all aspects of *Heads Up Hockey*, as well as beginning the integration of the instructional components into the game. During this time, we discovered that the targeted hardware platform of a 700 MHz processor and 256 MB of RAM was insufficient to handle 30 frames per second during game play conditions. The high resolution (800x600) of the graphics resulted in large amounts of data loaded in and out of memory. This necessitated a reduction in the game's resolution from 800x600 to 640x480 and the raising of the base memory requirements to 512 MB. These moves allowed the game to run at an acceptable speed of 30 frames per second on a computer that had at least a 700 MHz processor.

The first cutscenes were delivered in late December 2003 when the focus of development was on the between periods user interface, graphical elements and designing the karma system. Fine tuning of the gameplay and user interface adjustments continued through the beginning of 2004, but the focus of development shifted to integrating the cutscenes into *Heads Up Hockey*.[4] We ran an evaluation study to determine the instructional effectiveness of the cutscenes. As we predicted, the participants who saw a particular cutscene scored significantly higher on a questionnaire than the control group who saw no cutscenes. Further discussion and details of the evaluation can be found in Ciavarro *et al.* (2005).

On March 1st, 2004, the *Heads Up Hockey* development team began to move into a testing and refining stage. The concept of karma still required some refinement for which I would be primarily responsible. Various outstanding issues continued to emerge, however, which delayed serious effort towards

[4] It should be noted that the cutscenes were disabled for the purposes of the study in order to obtain a clearer evaluation of implicit teaching mechanism.

designing and implementing a karma system until July 2004. Our subject matter experts in Kinesiology came up with a list of activities that should result in negative and positive karma changes. It was my job to design a translation between a composite karma score and on-ice effects. I devised an operationalization of 'karma' that increases the likelihood of losing for those players that play aggressively or negligently. Since competitive game players do not like to lose, I hoped they would associate losing with negative behaviour. If this connection was made then I expected that the players would positively change their behaviour in order to win.

Similar to the summer of 2003, we wanted to test *Heads Up Hockey* at hockey camps in August 2004. The goal was to have the cutscenes working alongside the karma system. The cutscenes, AI and administration screens continued to be refined, but the karma system had only recently been designed. It was up to me to implement karma and have all other outstanding issues completed before the hockey summer camp. We made the deadline and I made the trip to the summer camp to observe *Heads Up Hockey* in action. Again, there was no formal measurements made of the participants' performance, but the post-session dialogue with participants reaffirmed our view that the game was ready to undergo systematic evaluation.

With the 2003 summer camp experience in mind, I set about cleaning up all remaining software bugs within the gameplay, increasing the customizability of the game settings to accommodate multiple game configurations for the purposes of testing, and adding a few tweaks to the user interface that would improve the overall usability of the game. To help achieve these objectives, we hired a dedicated game tester to play the game one day per week.

Further improvements to *Heads Up Hockey* through the fall of 2004 focused on trying to move the game to a web distribution model. Our reasoning for exploring this route was that we needed participants to play the game in order to evaluate the learning design. The web is a powerful draw for such a purpose because of the number of people it can reach. In order for this to work, we needed to devise a way to upload data, as well as keep the game up to date. In the end, despite a system that was fairly robust, we decided to go with an offline distribution model because it afforded us more control over the players' gaming experience.

In January 2005, there was another brief development push to accommodate my research objectives. One example of a change that needed to be made was the ability to separate the karma system from the cutscenes, since my study was to look at the effects of the karma system only. Details on the study itself can be found in the following chapters.

By summer 2005, all significant development on *Heads Up Hockey* had ceased. This project produced at least 3 distinct versions of the game that satisfied various research goals, such as examining the effectiveness of the cutscenes and the effect of "how to play" instructions on the level of game competency. The rest of this chapter will focus on the detailed mechanisms of the karma system, as well as give the reader a clearer picture of what playing

Heads Up Hockey is like. Design details on other aspects of the game (e.g. likelihood of injury calculation) will be presented where necessary.

4.3 Design elements of *Heads Up Hockey*

Heads Up Hockey was designed and developed without the use of proprietary tool kits or open source libraries, so there are many design points of interest, including: the physics system, the AI representation of team-mate and opponent behaviours, the multimedia processing of digital animation libraries and the Flash/Director interface. For the purposes of this thesis I will focus on the design of the karma system and its algorithms, which is intended to reduce unsafe hockey gameplay behaviours. Understanding the karma system's support processes will flesh out the scope of our instructional method. The support processes include: 1) injury calculations, 2) influencing player attributes and 3) line editing[5].

We identified two types of negative behaviours that we wanted to reduce: negligent and aggressive. If the participant does not deal with the injured player by permanently removing them from the game lineup, we consider this an act of negligent behaviour and we penalize the participant's composite karma level accordingly. Injuries can occur whenever a player is hit by an opposing player. The other component to negative karma is aggression, which was defined by selecting a set of in-game transgressions, such as charging or hitting the goaltender, and assigning karma penalties to them. An important difference

[5] "Line editing" refers to the assignment and re-assignment of the player roster to each of the three lines

between negligent and aggressive karma penalties is that aggressive penalties are one time subtractions (per transgression), whereas negligent karma penalties are continually accrued until the negligent behaviour ceases. The reason why we chose to have negligent karma continually accrue as opposed to a one time charge is because of the nature of neglecting a player. Negligence is an on-going phenomenon that does not take place at a single time event, whereas an aggressive act is a singular occurrence that befits a one time karma deduction.

The goal of the karma system is to change a player's behaviour to be less aggressive and negligent through a reward/penalty system that is based on specific acts. Composite karma, as we define it, is an integer between -100 and +100 whose value has corresponding effects on a player's ability to win, as outlined in Table 4.1. As mentioned above, there are two types of negative karma that can be incurred by a player: aggressive and negligent. We defined several aggressive acts that will cause an accumulation of aggressive karma. In addition, negligent behaviour, as defined above, will result in a periodic reduction in a player's composite karma score. A player gains positive karma by refraining from our chosen set of aggressive and negligent behaviours. The maximum rate at which a player could gain karma was 1 point every two seconds. This makes it easier to achieve large gains in negative karma because a participant can actively exhibit aggressive behaviour, whereas the only way to accrue positive karma is to play responsibly and wait for the positive accumulations to have effect.

Table 4.1: Aggressive and negligent behaviour penalties

Aggressive	**Karma Penalty**
Interference	-7
Run the Goalie	-11
Multiple Hits	-11
Charging	-11
Hit multiplier	Every penalty multiplied by # of successive hits
Negligent	**Karma Penalty**
1 injured player in lineup	-0.2/sec
1 injured player on ice	-1/sec
> 1 injured player in lineup	-1/sec
> 1 injured player on ice	-1/sec

Values taken from the summer 2005 testing sessions

It is important to note that the karma system was designed to be tightly bound, in that aggressive, negligent and positive karma can affect the composite karma score simultaneously. Furthermore, the aggressive penalty algorithm compliments the negligent penalty algorithm, and both of those work with the reward framework to alter the composite karma score. This gives rise to the problem of observing which class of behaviour a player is exhibiting more and whether it is changing. To assist with this problem, we keep track of the accumulations of negligent and aggressive karma, which give insight into what specific behaviour was being exhibited. However, the teaching mechanism is based upon the interaction between negative and positive karma, which means the evaluation of the teaching mechanism cannot only look at the separated accumulations of negative karma.

The metric that best measures overall behaviour (aggressive, negligent and positive) is the composite karma score. We cannot add the aggressive karma accumulations with the negligent totals to get a composite behaviour score because they represent two distinct types of negative karma. For example, if a player accrues 50 negative karma points by leaving an injured player in the lineup, the player will not accrue any positive karma during the negligent act. On the other hand, if a player accrues 50 negative karma points by playing aggressively, that player can still accumulate positive karma points. Thus, even though both players have accumulated 50 negative karma points, the aggressive player would have had a chance to accrue positive karma. This phenomenon is not captured by separately examining each class of negative behaviour. Therefore, the composite karma score is the best metric to evaluate the overall effectiveness of the teaching mechanism, whereas the two negative karma accumulators can best give insight as to which kind of behaviour the game most affected.

4.3.1 Influencing player attributes

Now that we have defined the function for how composite karma increases and decreases, we now need to define the effect that *karma*[6] has on gameplay. It is important that high *karma* scores lead to winning and low *karma* scores lead to losing, irrespective of individual game playing ability. If the karma system can easily be circumvented by superior game playing skills then we may no longer be associating game success with positive behaviour. The basis of our teaching mechanism relies on the player associating positive behaviour with winning. Therefore each limit of the *karma* scale needs to significantly impede or assist the game player. We achieve this by adding *karma* as a parameter to the

[6] The rest of this document will use an italicized *karma* in place of 'composite karma'

effective player attribute function. For example, the function for how hard a player can shoot the puck is as follows:

Code Sequence 4.1

```
return injuryEffects (((rink.energyeffectiveness + energy) /
float(rink.energyeffectiveness + 100)) * shotpower + karmadampen)
```

Every player attribute, such as `shotpower`, `speed` or `accuracy`, is a number between 0 and 100 (a full listing of player attributes can be found in Appendix E). The baseline shot power ability of this player is represented by the `shotpower` variable. We will first look at the role `karmadampen` has on this value, which has a 'dampening' effect on the team's *karma* score for the purpose of addition to player attributes. The reason we dampen the *karma* score is because it has a minimum value of -100. Without dampening, a player might not be able to shoot the puck because shot velocity is based on the `shotpower` variable, which has a maximum value of 100. The dampening factor for our testing was set at 6 (i.e. the raw *karma* score was divided by 6 before being applied to any player attributes). Thus, *karma* alone could increase or decrease any attribute by a maximum of 16 points.

Karma also has secondary influences on a player's ability, such as in our shot power example. Note that the player's energy level is a multiplicative factor in the effective value of the attribute. The `energyeffectiveness` modifier changes how much power the energy level has on a player. For this study, the effect of energy was made to vary between 0.75 and 1. *Karma* is involved at this stage by affecting the endurance of the player. In other words, if a player has a low *karma* score, they will lose energy faster which in turn lowers their abilities, on top of the direct effects of *karma*. It is important to note that a player has no way of viewing the effective levels of any of their player's abilities. The player can only view the composite energy levels of each line and baseline attributes of individual players.

The final modifier of a baseline attribute from Code Sequence 4.1 is the `injuryEffects` function. This function severely limits a player's ability in the case that a player has incurred a brain injury. In fact, the `injuryEffects` function can decrease any attribute by up to 33%, if a player has incurred a full 100 score of brain injury points. An extra effect of incurring an injury is that participant control of the injured player is disrupted thereby mimicking the disorientation a concussed player might exhibit. *Karma* also has an effect on the `injuryEffects` function by influencing the likelihood of a player being injured, thereby altering the chances that this function will come into play.

4.3.2 Injury occurrence calculation

The injury occurrence mechanism is an important component of the karma system. An injury could occur on every hit, but there are many factors (including *karma*) that affect the chances of an injury occurring. However, at the highest level, the actual chance of injury calculation is quite simple.

Code Sequence 4.2

```
coi = 1000 * (injurypoints - mininjurypoints) /
(maxinjurypoints.float - mininjurypoints.float)
```

The chance of injury (`coi`) is a number between 0 and 1000, based on the variable `injurypoints`. A random number is then chosen from a uniform distribution between 0 and 1000, and if that number is less than the value of `coi` we say that an injury has occurred. Thus, the higher the `coi` value, the greater chance that a particular hit will cause an injury. The derivation of `injurypoints` is the main value in this calculation, and also where the complexity arises. The following is the snippet of code that determines the value of `injurypoints` for every on ice collision. For the purposes of the injury calculation, we defined a hit as a one way collision between a hitter and a hittee.

Code Sequence 4.3

```
injurypoints = hitlevel * injurycalcfactors.hf[1] +
      sprite(hittee).injury_state * injurycalcfactors.isf[1] -
      sprite(hittee).injury_resistance * injurycalcfactors.rf[1] -
      k * injurycalcfactors.kf[1] + ht * injurycalcfactors.htf[1] +
      sprite(hitter).getAttribute("aggressiveness", "af") *
      injurycalcfactors.af[1]
```

There are six baseline parameters and six modifiers that determine the final value of `injurypoints`. The variable `hitlevel` is a discrete number with a maximum of 3 that represents how much force was involved in the collision. The hittee's current `injury_state` (concussed or not) and overall resistance to injury play a part in determining the likelihood of injury. *Karma*, the type of hit (boarding, open ice, etc.) and the hitter's level of aggressiveness also have a direct role in whether an injury occurs. The six modifiers to these baseline variables allowed us to easily tweak the influence of each of these values in order to control the average number of concussions that occurred in each game.

4.3.3 Line editing

Taking injured players out of the lineup is crucial to winning at *Heads Up Hockey* because of the negative accumulations of karma that occur if injured players are ignored. The design and implementation of the line editing mechanism was complex since it required the Director game engine to converse with the Flash user interface. Some of the issues were: populating the interface with the proper players, sending changes back to the main game engine, designing an intuitive method for altering the lineup and ensuring a smooth transition to and from the line editor and gameplay. In order to accomplish these goals, we used the built-in capabilities of Flash and Director to swap simple commands and data structures with each other.

The design of the line editing user interface and its underlying algorithms was cumbersome. There were a surprising amount of cases we needed to handle. The structure of the problem was as follows: there are nine empty slots divided into three lines, a player may not appear twice on the same line, but can appear on other lines, and all lines must have 3 players. Due to complications from our attempts at improving intuitiveness (e.g. dragging a player to another line will remove the target player and double shift the dragged player) it took several iterations to achieve a working solution. The most problematic issue in getting line editing fully operational was the ability to fully pause the game and enter the line editing screen. The game was not originally designed with pause in mind, so we had to retrofit this feature into the game. Everything from player positions, to timers, to camera positions needed to be saved for later reinstatement.

Figure 4.3: Pre-game lineup editor

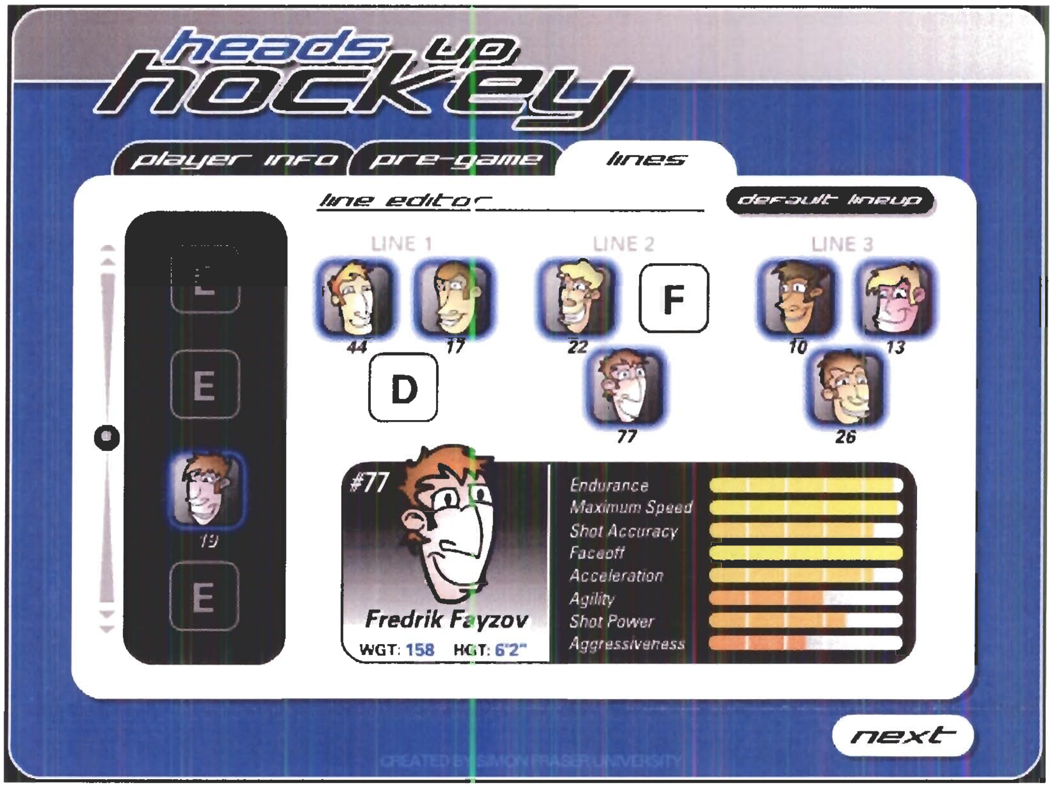

4.4 Gameplay design

This section provides some textual and pictorial examples of gameplay to help describe the experience of playing *Heads Up Hockey*. At the start of every game, the participant is shown a main menu, which is followed by a series of instructions. The instructions describe how to play the game and some of the encouraged strategies, including how to remove injured players. Basic skills such as passing, shooting and skating are covered. More advanced skills, such as one-timer shots, are covered as well.

Figure 4.4: Heads Up Hockey main title screen

Settings and Quit were disabled during our testing

After the instructional slides, the player has the opportunity to select their lines (Figure 4.3). There is a database of 12 players to fill 9 possible slots. Each player has a different name, face and skill set. Once a player sets a lineup, it is stored for future use by that player, which saves the effort of trying to reform lines before every game. The in-game line editor is similar to the before-game line editor with the major difference being that, during gameplay, participants are limited to a nine player roster with double shifting permitted. If a player has sustained a brain injury, the in-game line editor will indicate this.

Figure 4.5: Sample instruction slide from Heads Up Hockey

After setting the lines, the player is sent right to the opening faceoff. Of course, the actual gameplay is where a lot of development effort was spent in terms of physics, graphics and controls. The action is quick and fluid, which challenges participants to keep pace immediately from the opening face-off. The fast pace of the game is compounded by the fact that it is 3 on 3 hockey, which means there is plenty of ice for skating. We emulated the *de facto* game controller setup from other hockey games in order to reduce the learning curve for experienced players. In addition, we kept the number of buttons to a minimum: one button for pass and change player, one button for shoot and speed burst, the analog stick for skating, and two shoulder buttons for line changes.

Figure 4.6: Screenshot of Heads Up Hockey gameplay

It is amongst this fast paced gameplay, colourful graphics and fluid animations that the karma system operates. The main benefit of our karmic teaching mechanism is that it is designed to avoid disrupting the player's engagement with the game. Individual karma events are not overtly discernible, having only a cumulative effect. One result of this is that a player is unlikely to acquire the desired behaviours with only a handful of plays of the game. This means that the player has to figure out on their own terms what they need to do in order to be successful at *Heads Up Hockey*. This is similar to "situated cognition" and "constructivist" learning environments (J. S. Brown *et al.*, 1989; Jonassen, 1999), in that we do not impart the "know what" but rather foster the "know how".

CHAPTER 5: STUDY DESIGN

5.1 Purpose

The goal of this study is to determine whether the karma mechanism built into *Heads Up Hockey* decreases the aggressive and negligent gameplay behaviours of participants. We do not expect all participants to be able to recognize that their behaviours are changing, but we do expect most participants to have some sense of the game's purpose. Although we do not attempt to transfer any specific content (we are interested in behaviour change), the game has the potential to be a knowledge transfer tool through the use of the aforementioned cutscenes. Further discourse regarding this opportunity and how to take advantage of it appears in Chapter 7. The remainder of this chapter details the objectives and methods of the study. The results of the study are presented in Chapter 6.

5.2 Objectives

The primary aim of this study was to determine whether there was a significant positive change in mean *karma* (behaviour) within the experimental group versus the control group. To achieve this goal, the first step was to determine whether the karmic teaching mechanism operated as expected. That is, we needed to examine whether high *karma* scores resulted in a greater ability to win, and vice versa. Furthermore, we wished to examine the constituent components of negative karma (aggression and negligence) to determine whether these behaviours were reduced over time and by what rate within the primary experimental group. Another goal of this study was to discover whether the experimental questionnaire respondents had a greater tendency to believe less aggression and negligence will lead to winning at *Heads Up Hockey*, as compared to the control group. Finally, we introduced a second experimental condition whose treatment and objectives are the inverse of the primary experimental condition in that these participants are expected to win by playing more aggressively and negligently.

5.3 Methods

5.3.1 Participants

Previous experience with testing *Heads Up Hockey* (*HUH*) revealed that the game remained engaging over a period of time only for certain age groups (Ciavarro *et al.*, 2005). From debriefing the test administrators and my own observations, it was clear that those participants who were 15 or over were not enthusiastic about playing the game for extended periods, perhaps because the game was not powered by the latest 3D graphics employed by today's popular digital games. In contrast, those participants aged 10 to14 years did not seem as distracted by the lack of a 3D graphics engine. This observation led us to restrict our study's participant age range to 10 to 14 years old in an attempt to maximize interest in the game.

We drew upon three separate subject pools. We had 16 participants from a lacrosse team, 29 participants from a summer camp and 29 participants recruited via a flyer, amounting to 74 participants. We received informed consent from each participant and, where applicable, their parent or guardian. For the purposes of a post-study interview, we received consent from the summer camp participants and their guardians. We did not conduct interviews with the other participants.

We were able to secure participation of the lacrosse team by promising them 500 dollars, while the Surrey summer camp class was recruited with the promise of amusement park passes, and finally individuals attending Burnaby summer camps were recruited via a flyer that promised $30 for their full participation.

5.3.2 Experimental groups

There were two experimental groups and one control group. The first experimental group we labelled *kep*, which stands for *K*arma *E*ffective *P*ositive. This group received the 'positive' treatment that rewards less aggressive and negligent behaviour with enhanced player attributes, and thus a greater likelihood of scoring and winning. The second experimental group we labelled *ken*, or *K*arma *E*ffective *N*egative. This condition was opposite to that of the *kep* group in that aggressive and negligent behaviour was rewarded with enhanced player attributes and a greater likelihood of winning. The control group we labelled *k0p*, or *K*arma *0* (not effective) *P*ositive. The control group did not experience *karma* effects (i.e. *karma* effectively remained at 0) but we continued to record the control group's actual *karma* level because *karma* was not only an influencer on gameplay under experimental conditions, but it was also our gauge of behaviour.

Thus, *karma* as a behaviour indicator was recorded for each group, experimentals and control, to allow for comparisons between groups.

Participants were randomly assigned to groups and were not informed about their group identity. However, to limit the possibility of bad advice being passed between participants, we did instruct talkative players that each game had a different setting and what works for one player might not work for another. Table 5.1 shows the breakdown among participant pools.

Table 5.1: Break down of participants across subject pools and condition

	Lacrosse	**Summer Camp**	**Recruits**	Aggregate
kep	6	11	11	**28**
k0p	5	9	10	**24**
ken	5	9	8	**22**
Sum	16	29	29	**74**
Age x	10.7	13.2	12.0	**12.2**
Age σ	0.7	1.3	1.6	**1.3**

5.3.3 Test session overview

Each player participated in three separate sessions of 5 games for a total of 15 games each. After a participant completed their 15th game, they were administered a 10-item questionnaire. There were no instructions from us other than to have fun. There were between one and three study administrators (depending on session size) present to answer questions if they arose. Each session was scheduled at least one, but no more than six, days apart. We required at least one day in between sessions to enhance the likelihood of meaningful reflection and to reduce the chance of “burnout” playing *HUH.* Burnout is defined here as overexposure to a game that causes disinterest and drop in attention to gameplay. We would have liked to cap the maximum reflection period at less than six days, but this was compromised to work with the participants’ schedules. The average reflection period for the lacrosse team was six days, whereas the summer camp and flyer-recruited participants had an average of three days between sessions.

The lacrosse team and individual recruits played on laptops that were available through the Motor Behaviour Laboratory at SFU. The summer camp participants played on desktop machines, which were able to run *Heads Up Hockey* at a marginally faster frame rate. For both the lacrosse team and the flyer-recruited participants each session had a different number of participants. This was due to not being able to collate everyone’s schedule into a single time slot. The average group size was around four participants for these subject pools. We were usually able to avoid having a participant play the game alone.

5.3.4 Test session detail

At the first session, participants would be presented with a screen prompting them for their first and last name. Next, they would be asked for their age and gender. After we had collected this information, the player moved to the game’s title screen. Before each of the 15 games, the participant was presented with 8 slides about how to play *Heads Up Hockey.* There was no requirement for how long a participant must read each slide, but anecdotal observation revealed that players spent several minutes during the first session absorbing the instructions. The slides were generally disregarded and quickly skipped through during later games.

After the slides, the player had the ability to form three lines from a roster of 12 players, each with slightly different attributes. After the player chose their lines, they were saved so that the line selection process did not need to be repeated before every game. With the lines set, the player began actual gameplay. Each period consisted of two minutes of stop-time play. On average, a single game took just over 10 minutes, which meant each test session lasted just under an hour. During the game, we recorded aggressive behaviours such as hard checking, running the goaltender and hitting a player who does not have the puck. We also recorded negligent behaviours such as leaving an injured player in the line-up, and putting a line on the ice that has a brain injured player (Table 4.1).

The session was complete once the participant finished his or her 5th game. As an added motivational tool, cumulative statistics were displayed at the end of the games for players to review. For example, the number of wins and losses, goals for and against, and total points are presented after each game. When the player returned for their 2nd and 3rd sessions, their lines and cumulative statistics picked up where the player left off. Upon the completion of the 3rd session (15th game), the player was asked to fill out a questionnaire that gathered information on the attitudes they recalled having during their *Heads Up Hockey* experience.

5.3.5 Data collection

During gameplay, *Heads Up Hockey* records a host of information. Regular hockey statistics are compiled, such as shots, goals and attack zone time. In addition, a play-by-play file is recorded which contains every incident of aggressive behaviour alongside regular hockey plays (e.g. a shot or a goal). Sample data and a play-by-play file are presented in Appendix A and C. We also recorded the mean *karma* levels on a per period basis. *Karma* is sampled once every second and added to the accumulator *kacc* (1). To get a per-period mean *karma* score (*ppavgk[per]*), the value of the accumulator is subtracted by the sum of the previous period's accumulator values, and divided by the number of seconds in a period (2). The mean *karma* score for a game (*gavgk*), which we use as our base unit of analysis, is the average of the three per period *karma* means, or the final accumulator value divided by the number of seconds in a game (3).

(1) $kacc = kacc + karma$, once per second

(2) $$ppavgk[per] = \frac{kacc}{mpp*60+spp} - \sum_{i=1}^{per-1} ppavgk[i]$$

(3) $$gavgk = \frac{\sum_{i=1}^{3} ppavgk[i]}{3} \text{ or } gavgk = \frac{kacc}{3(mpp*60+spp)}$$

In addition, the two kinds of negative behaviours, aggressive and negligent, were assigned their own accumulators. These values were not averaged over time, rather they simply accrued value as negligent and aggressive karma was incurred. There is no maximum limit on the aggressive and negligent karma accumulators. All karma measures were initialized to zero at the start of every game.

It is important to note the difference between the way composite karma and negative karma are recorded. Composite karma is a combined measure of negative and positive behaviours, whereas the negative accumulators only capture negative behaviours. This distinction is important because the evaluation of the implicit teaching mechanism must take into account more than just a simple accumulation of negative behaviour. A player's average behaviour over the entire game is what is most important. To further illustrate this point, if a player exhibits generally positive behaviour for the first 90% of the game and is well ahead on the scoreboard, there is a natural tendency to relax and begin to drift away from successful habits. It is possible then that a player that has played the entire game in a positive fashion to accumulate a large amount of negative karma in the last 10% of the game. Without examining the mean behaviour score (*karma*), we would miss the generally positive behaviour that the player exhibited throughout the game.

Once the 15th game was complete, we administered a 10-item questionnaire. The goal of the questionnaire was to determine whether the participant was able to correctly reflect on and identify the behaviours that the game was trying to instil. The reason why we wanted to look at this was to establish a measure of each participant's level of awareness about their game playing behaviour and the reasons why it changed. The idea was that if a player was aware that their behaviour had changed, then they would be more likely to transfer that learned behaviour to other contexts, such as on-ice hockey games. A copy of the questionnaire is presented in Appendix B. For the summer camp pool, we requested permission to perform an audio interview. The goal was not to interview every participant, but rather those who seemed to exhibit the expected behaviours within the *kep* group (that is, increased their mean *karma* scores over the three sessions). This would allow us a deeper understanding of some of the factors that led this participant to the realization of the winning strategy. The interview itself was unstructured, but we did use a rough outline of questions, as presented below.

- "Can you remember around which game (1 to 15) you figured out how to win at *HUH*?"
- "What do you think was the most important thing to winning at *HUH*?"
- "Do you know how you came to this strategy? If yes, how?"
- "Is the way you won at *HUH* different from other hockey video games you have played? How so?"
- "Do you think *HUH* was trying to teach you anything? If so, what?"

5.4 Predictions

The goal of this study was to determine whether the karmic mechanism was effective at curbing aggressive and negligent gameplay behaviours in participants. To achieve our goal, the first step was to determine whether the karmic mechanism positively correlated *karma* to winning (goal differential) for the *kep* group and negatively correlated *karma* to goal differential for the *ken* group. The control group, *k0p*, should have no correlation between winning and *karma*. If the karma system influences behaviour as expected, we anticipate a significant difference, both in magnitude and positive rate of change of composite karma, between the *kep* and *k0p* groups. In addition, the separate accumulations of negligent and aggressive karma should decrease more rapidly over time in the *kep* condition versus the control group. The *ken* group should exhibit a slight increase in negligent and aggressive karma accumulations over the control group.

We have several secondary expectations as well. First, we expect the *ken* group to experience little change in *karma* over the three sessions. Since the *ken* group rewards least effort behaviours (it is easier to leave a concussed player in the lineup), we should see *karma* values slightly decrease (Zipf, 1949). The reason why we only expect a slight decrease in mean *karma* values within the *ken* group is because we anticipate participants to start off with a low behaviour score, thus not leaving much room for decrease. Also, we would expect that the *kep* questionnaire respondents will have a greater tendency to believe less aggression and negligence will lead to winning at *Heads Up Hockey*, as compared to both the *ken* and *k0p* groups. The *ken* group should believe most strongly that aggression and negligence lead to winning at *Heads Up Hockey*.

CHAPTER 6: ANALYSIS AND RESULTS

To begin the analysis, the quality of the data required scrutiny. Through observation, I discovered some participants would need to be excluded from the study because of disinterest or tampering with the game. To gauge disinterest, there was a question on the questionnaire that asked if a participant tried their best to win at *Heads Up Hockey*. If a person reported that they did not try their best and were also on our list of people observed to be exhibiting disinterest, they were excluded from the analysis. After checking our list of disinterested participants, we found they matched perfectly and three sets of data were discarded. With respect to tampering, there was a loophole in the game that allowed the participant to inadvertently reset the game settings, which resulted in the removal of another three participants from analysis. There was also a case of a participant who had previous experience with *Heads Up Hockey* and therefore was ruled ineligible after they completed their final testing session. In addition, some participants did not complete every game, or they tampered with the game in the 2nd or 3rd session. In these cases, we were left with less than a full complement of 15 *karma* data points for some participants. This data was excluded from analysis due to the repeated measures design.

Table 6.1: Summary of number of analyzed participants and games played

	Recruited	Analyzed	Games
kep	28	21	315
k0p	24	23	334
ken	22	21	309
Total	**74**	**65**	**958**

6.1 Analysis of the karmic teaching mechanism

I wanted to assess whether exposure to the karmic mechanism resulted in the expected outcomes. In other words, I aimed to determine whether the karma system was influencing players as expected by rewarding safe play with wins and negative behaviour with losses. For the *kep* condition, if the participants play with low aggression and negligence they should be rewarded with wins (i.e. positive goal differential). The *ken* condition should reward aggressive and negligent behaviour, thus *karma* should be negatively correlated with goal differential (i.e. low *karma* results in positive goal differential). Finally, the control condition *k0p* should have no correlation between *karma* and goal differential. The following three Figures (6.1, 6.2 and 6.3) show the scatter plots of mean *karma* scores and goal differential (home score minus away score) for each game within a single condition.

Figure 6.1: Scatter plots of *karma* and goal differential for kep

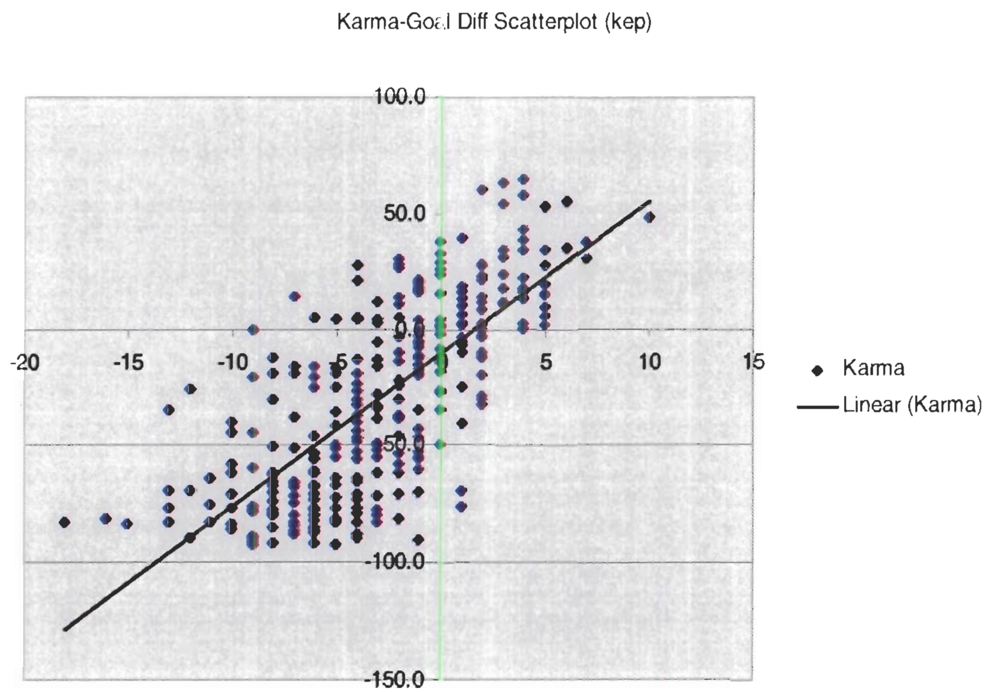

Figure 6.2: Scatter plots of *karma* and goal differential for ken

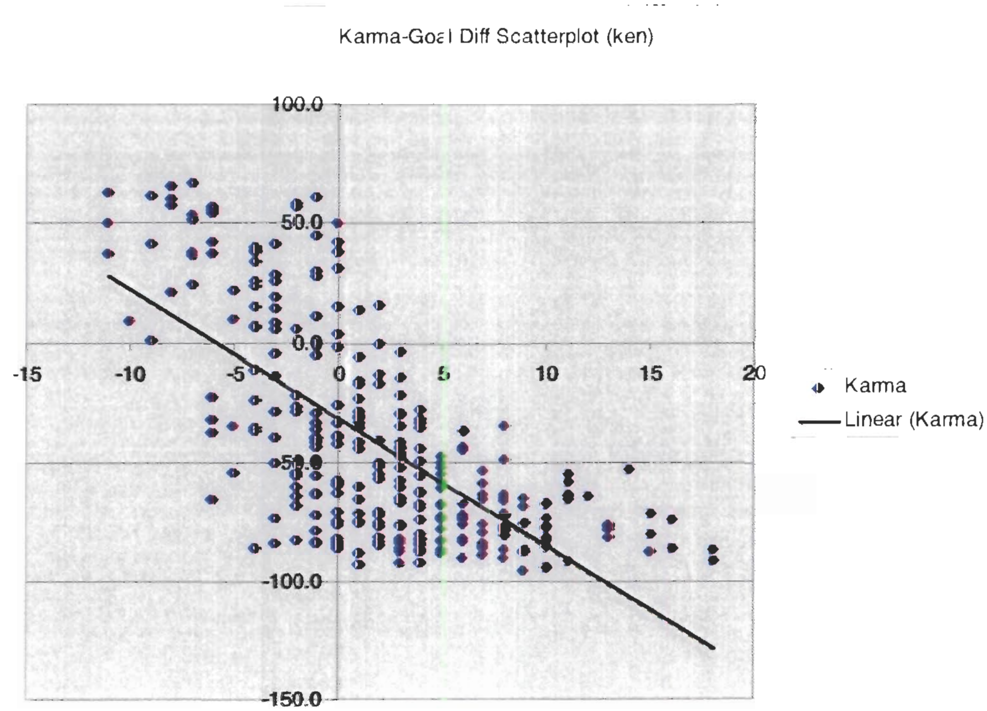

Figure 6.3: Scatter plots of *karma* and goal differential for k0p

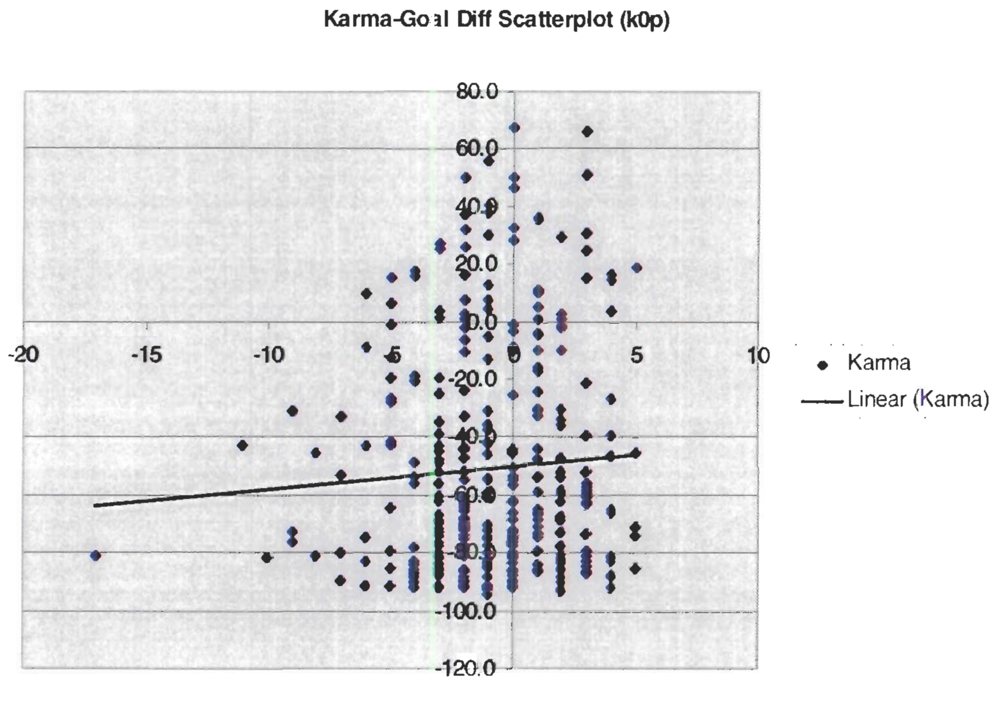

To help describe the relationship between *karma* and goal differential, I used Pearson's correlation on pairs of mean *karma* scores and goal differential for each condition (see Table 6.2). There were 315 pairs (i.e. games) for the *kep* condition, 309 pairs for the *ken* condition and 334 for the control group. The analysis showed *karma* and goal differential to be highly correlated under the *kep* and *ken* conditions. The control condition showed no correlation between *karma* and goal differential. This is an important result because if the control condition was rewarding either positive or negative behaviours, it would be acting like the other conditions which would make it difficult to compare results.

Table 6.2: Correlation between *karma* and goal differential

Condition	*n*	Pearson's *r*
kep	315	0.733
ken	309	-0.681
k0p	334	0.063

Both experimental condition r values are significant ($p < 0.0005$)

To further describe the relationship between *karma* and goal differential, I performed a linear regression analysis. The results of this analysis indicate that there is a linear equation between goal differential and mean *karma* under the *kep* condition with statistically significant coefficients ($y_{gd} = 0.082x_{karma} - 0.528$). In addition, there is no significant slope constant for the *k0p* group, indicating that there is no linear relationship between goal differential and *karma*. Finally, there is a linear equation for the *ken* condition that has a significant intercept and slope ($y_{gd} = -0.086x_{karma} - 1.499$). This indicates that *karma* is a predictor of goal differential under both experimental conditions, but is not a factor under the control condition. However, the adjusted R square for the *kep* and *ken* groups was 0.536 and 0.462, respectively. Intuitively, this means that there are other factors at work that determine goal differential, which is expected because *Heads Up Hockey* is a game of probabilistic events that can effect both *karma* and goal differential. Finally, the adjusted R square for the control group is almost zero (0.001) which indicates that there was little relationship between *karma* and goal differential within the control condition.

Table 6.3: Results of linear regression analysis

Condition	X Variable	Sig.	Intercept	Sig.	Adjusted R Square
kep	0.082	$p < 0.0001$	-0.528	$p < 0.02$	0.536
ken	-0.086	$p < 0.0001$	-1.499	$p < 0.0001$	0.462
k0p	0.005	$p > 0.05$	-0.606	$p < 0.03$	0.001

6.2 Analysis of the mean *karma* scores

The next step is to look for changes in behaviour across the experimental groups. We expect the control group (*k0p*) to exhibit essentially no change in *karma* scores (behaviour), whereas the experimental group *kep* should learn from the karmic mechanism and improve their behaviour. The *ken* group should exhibit a small negative change in behaviour. We only expect a small negative change for the *ken* group because we assume most players will naturally play with aggression and negligence, thereby leaving less room for significant drops in *karma* scores. The analysis started by breaking up the 15 *karma* scores per participant into three sessions. The 5 *karma* scores per session were averaged to give a single session score.

Participants were listed row-wise, each having three *karma* scores representing each of their three sessions. Each column was averaged to give a mean group score for that session. When plotted (Figure 6.4), these per session *karma* means clearly show a difference between the *kep* group and the others. As expected, the *kep* group increased their *karma* scores as time went along, whereas the control and *ken* groups did not appear to change their behaviour.

Figure 6.4: Mean *karma* per session plot

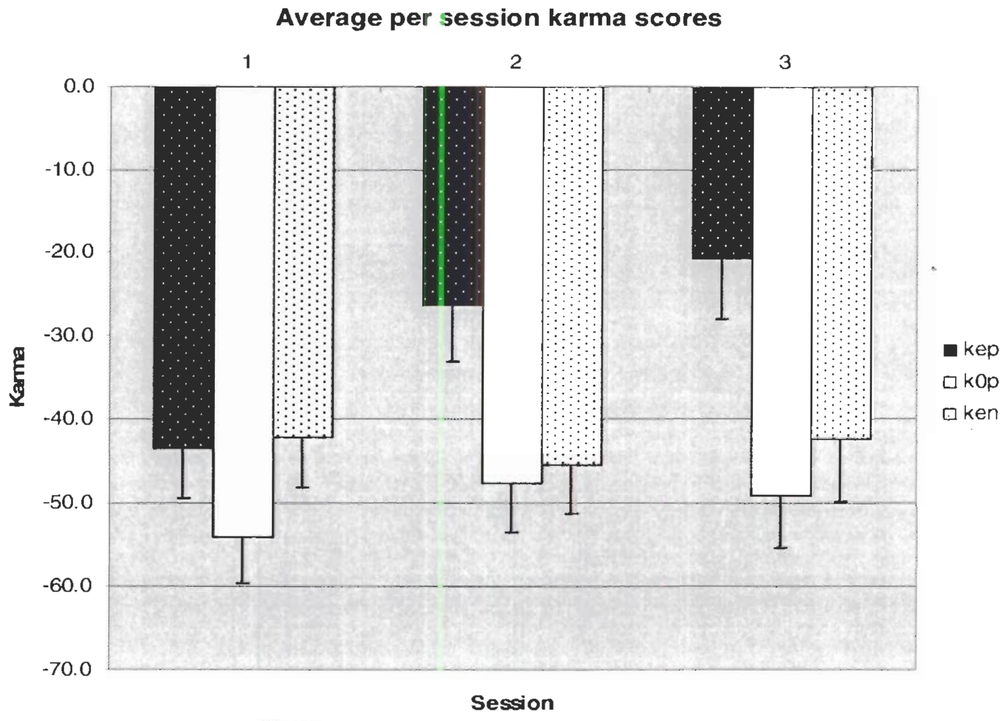

There was concern that the *ken* and *k0p* groups could be experiencing a floor effect. The minimum possible *karma* score was -100, which meant that when a participant reached the minimum, any further negative actions were essentially lost. By inspecting the distribution of the three groups of *karma* scores it becomes apparent that the *k0p* and *ken* groups are skewed to the left (Figures 6.6 and 6.7), which is an indication of a floor effect. In addition, the *k0p* and *ken* groups had 14% and 21% of their per period *karma* scores at the minimum threshold (-100), respectively. The *kep* group only had 3% of the per period *karma* scores at the minimum threshold. This is an indication the *kep* scores were distributed more evenly, thereby easing the likelihood of a significant floor effect (Figure 6.5).

Figure 6.5: Distribution of the *kep karma* scores (n = 315)

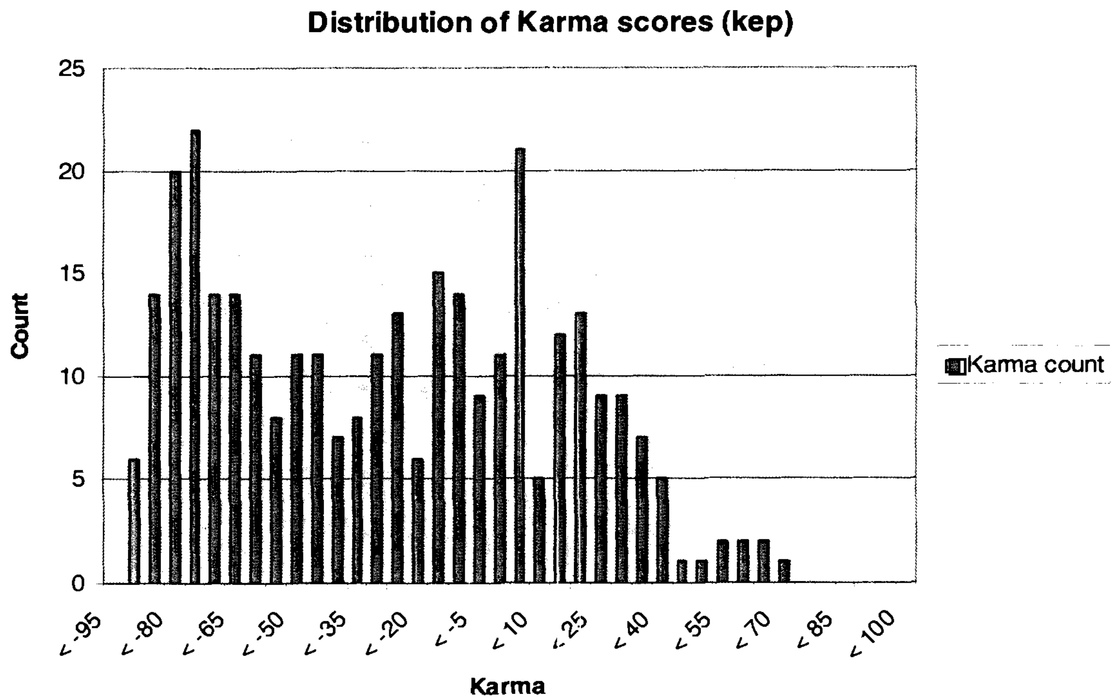

Figure 6.6: Distribution of the *k0p karma* scores (n = 334)

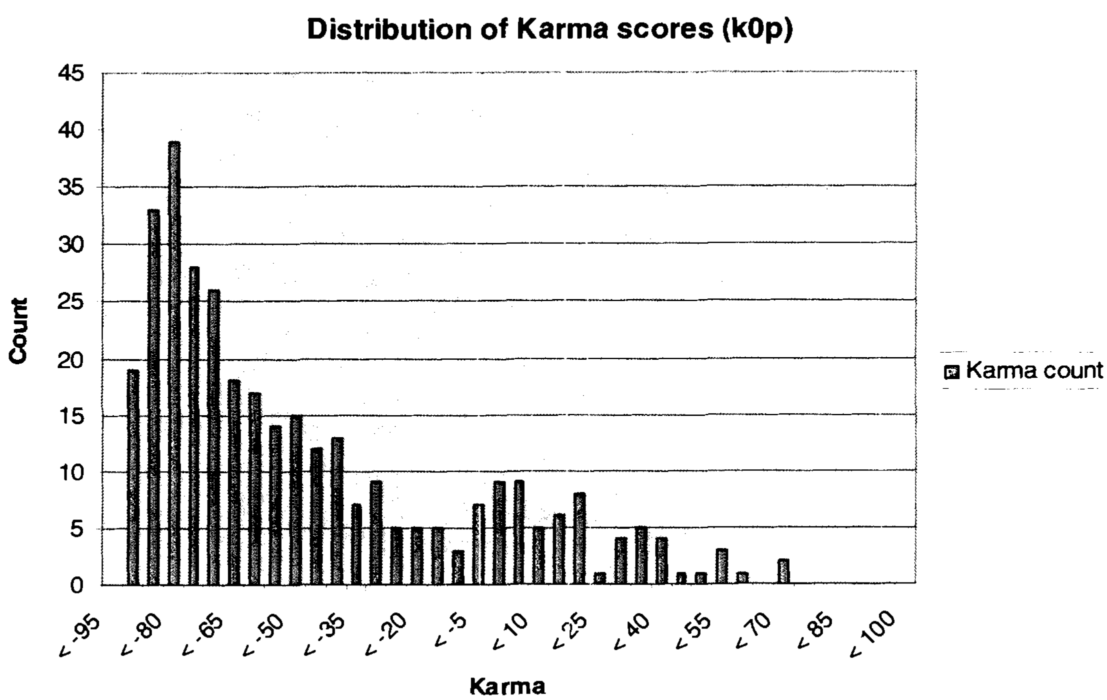

Figure 6.7: Distribution of the *ken karma* scores (n = 309)

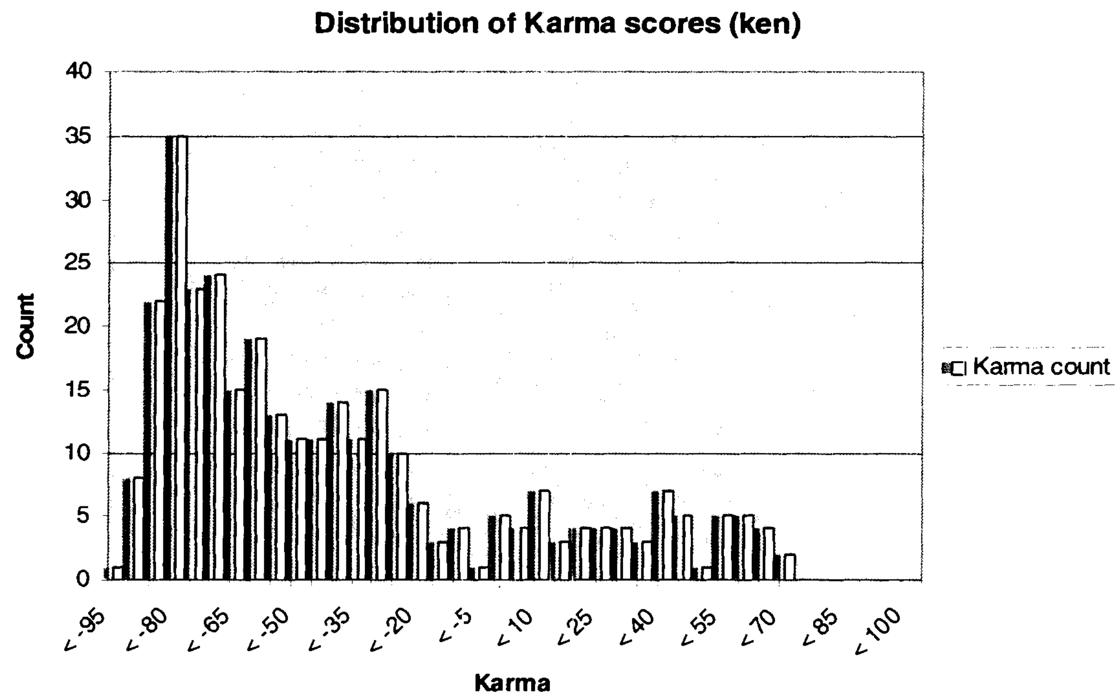

To show that the *kep* group significantly increased their *karma* scores, and thus positively changed their behaviour, I employed a two-factor repeated measures ANOVA in SPSS version 13. The two factors (condition and session) each have three levels (*kep, kOp* and *ken* for condition and 1st, 2nd and 3rd for session). This particular analysis does not allow incomplete data sets, thus one *ken* and two *kOp* participants had to be dropped. We found the main effect (session) to be significant ($p = 0.021$), indicating that time exposure alone had an effect on *karma.* Next, we found that the interaction between condition and session to be significant ($p = 0.026$), as summarized in Table 6.7. This indicates that certain conditions increased *karma* scores over time more so than others. The expectation was for the *kep* condition to increase their karma scores more so than the control group. Indeed, we can see in Figure 6.8 the considerable positive effect the *kep* group had in improving overall *karma* scores. The main session effect can be seen in the magnitude differences between *kep* and the other groups, while the session*condition interaction is apparent from observing the greater slope of the *kep* line versus other plots.

Table 6.4: Within-Subjects factors for overall *karma* data

Session	Dependent Variable
1	Session1
2	Session2
3	Session3

Table 6.5: Between-Subjects factors for overall *karma* data

		N
Condition	kOp	21
	ken	20
	kep	21

Table 6.6: Descriptive statistics for overall *karma* data

	Condition	Mean	Std. Deviation	N
Session1	kOp	-55.2048	26.23813	21
	ken	-40.5200	27.01074	20
	kep	-43.5238	27.19807	21
	Total	-46.5113	27.13595	62
Session2	kOp	-50.2810	28.45844	21
	ken	-44.2067	26.59929	20
	kep	-26.2952	31.30761	21
	Total	-40.1973	30.23387	62
Session3	kOp	-49.4952	28.82753	21
	ken	-42.0925	33.58615	20
	kep	-20.7333	33.26251	21
	Total	-37.3653	33.76069	62

Table 6.7: Tests of Within-Subjects effects on overall *karma*

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
session	Sphericity Assumed	2611.057	2	1305.528	3.984	.021
session * Condition	Sphericity Assumed	3750.868	4	937.717	2.862	.026
Error(session)	Sphericity Assumed	38664.206	118	327.663		

Table 6.8: Tests of Within-Subjects contrasts for overall *karma* data

Source	session	Type III Sum of Squares	df	Mean Square	F	Sig.
session	Linear	2496.212	1	2496.212	6.035	.017
	Quadratic	114.844	1	114.844	.475	.493
session * Condition	Linear	3227.664	2	1613.832	3.902	.026
	Quadratic	523.204	2	261.602	1.082	.345
Error(session)	Linear	24402.688	59	413.605		
	Quadratic	14261.518	59	241.721		

Measure: *karma*

Figure 6.8: Plot of mean *karma* over session

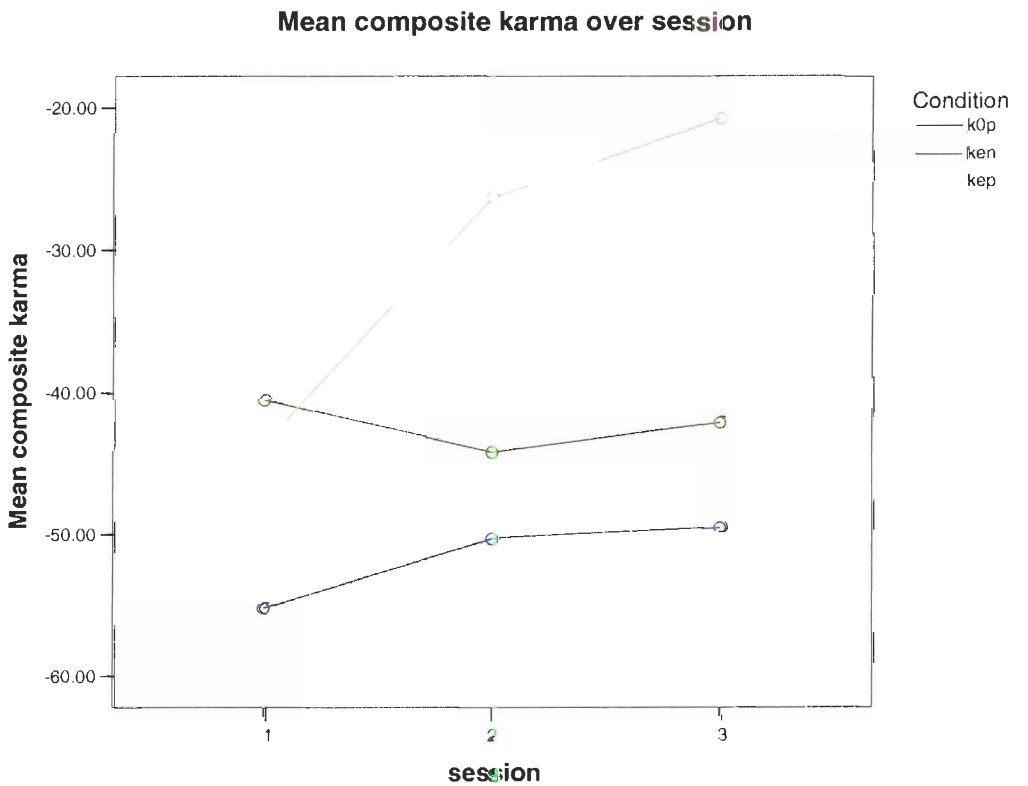

Generated from SPSS ver. 13

Given the significant condition*session interaction, we ran three separate 1-way ANOVAs with repeated measures on each condition in order to determine which groups significantly changed their *karma* scores over time. The *kep* condition was the only group that had a significant main effect (p = 0.0002), which, based on a positive trend analysis, revealed a significant linear (p = 0.0004) trend. The *k0p* group (p = 0.423) and *ken* group (p = 0.868) did not significantly change over time. These results are summarized in Table 6.9.

Table 6.9: Results of ANOVA with repeated measures across conditions for *karma*

Condition	Source	Type III Sum Of Squares	df	Mean Square	F	Sig.
kep	Session	5930.150	2	2965.075	10.589	.000
	Error	11200.517	40	280.013		
k0p	Session	402.219	2	201.110	.878	.423
	Error	9158.007	40	228.950		
ken	Session	136.893	2	68.447	.142	.868
	Error	18305.682	38	481.728		

Table 6.10: Tests of within-subjects contrasts for kep *karma*

Source	session	Type III Sum of Squares	df	Mean Square	F	Sig.
session	Linear	5453.761	1	5453.761	18.114	.000
	Quadratic	476.389	1	476.389	1.840	.190
Error(session)	Linear	6021.719	20	301.086		
	Quadratic	5178.798	20	258.940		

Measure: kep negligent karma

Figure 6.9: Plot of mean *kep karma* over session

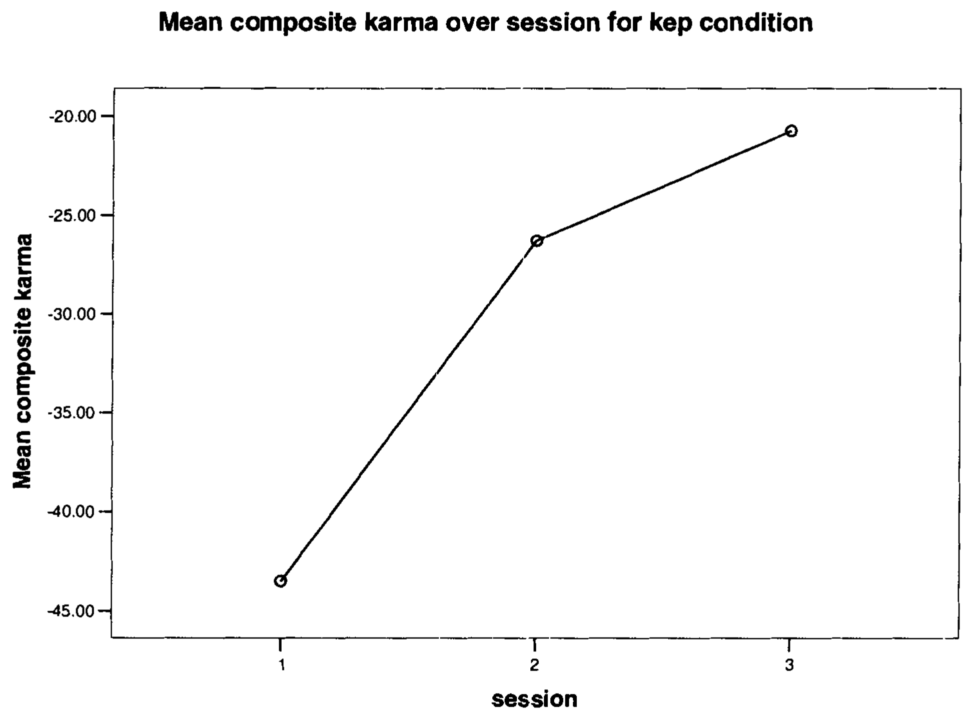

6.3 Analysis of aggressive and negligent karma

Negative karma consists of two types: aggressive and negligent. Each was recorded separately, and this section looks at how these constituents of negative karma changed over time and across conditions. This analysis will give us a better understanding of what was behind the significant results found in the last section. Similarly to the analysis of the composite karma scores, we expect negligent and aggressive karma to decrease over time within the *kep* group versus the control group. The *ken* group should slightly increase its negligent and aggressive totals as compared to the control group.

The method of analysis is the same as the analysis of mean *karma* scores, except the analyzed values are the average accumulations of negligent or aggressive karma. To begin, we plotted the mean aggressive and negligent karma scores over the three sessions. We can observe from these plots that negligent karma appears to be changing as expected, with the *kep* group reducing their negligent karma and the *ken* group exhibiting an increase (Figure 6.11). On the other hand, aggressive karma does not show much change between sessions for any condition (Figure 6.10). This suggests that the composite karma changes were being driven by the effects of negligent karma rather than aggressive karma. This could mean that players believed that editing their lines to remove injured players was more important to winning than being less aggressive. Further analysis is required to understand the influence of aggressive and negligent karma.

Figure 6.10: Mean aggressive karma per session across conditions

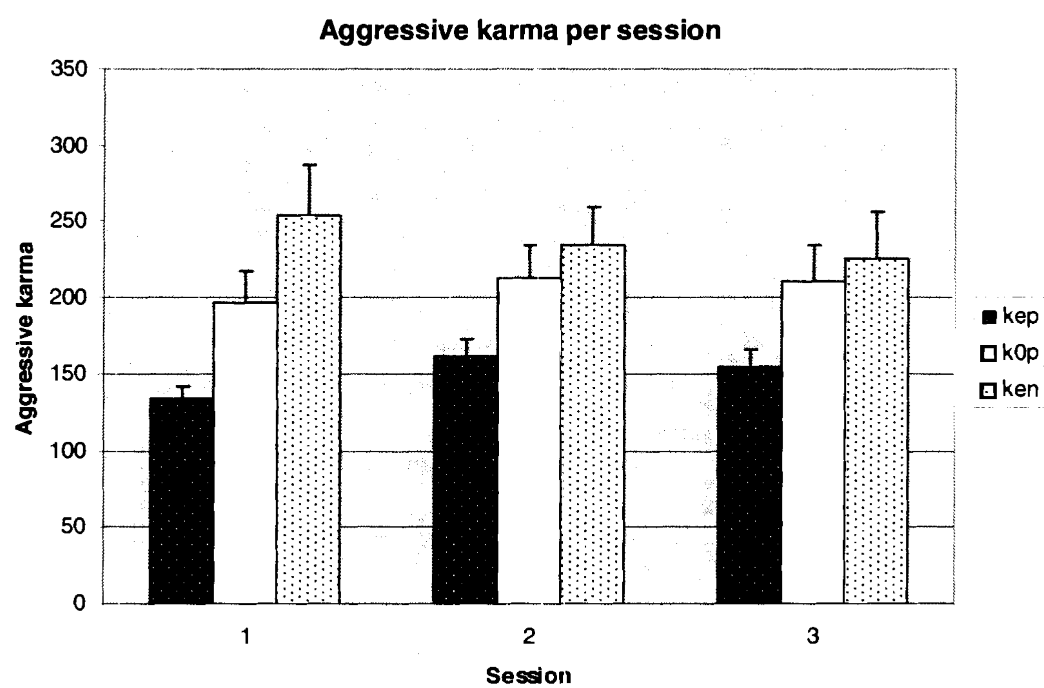

Figure 6.11: Mean negligent karma per session across conditions

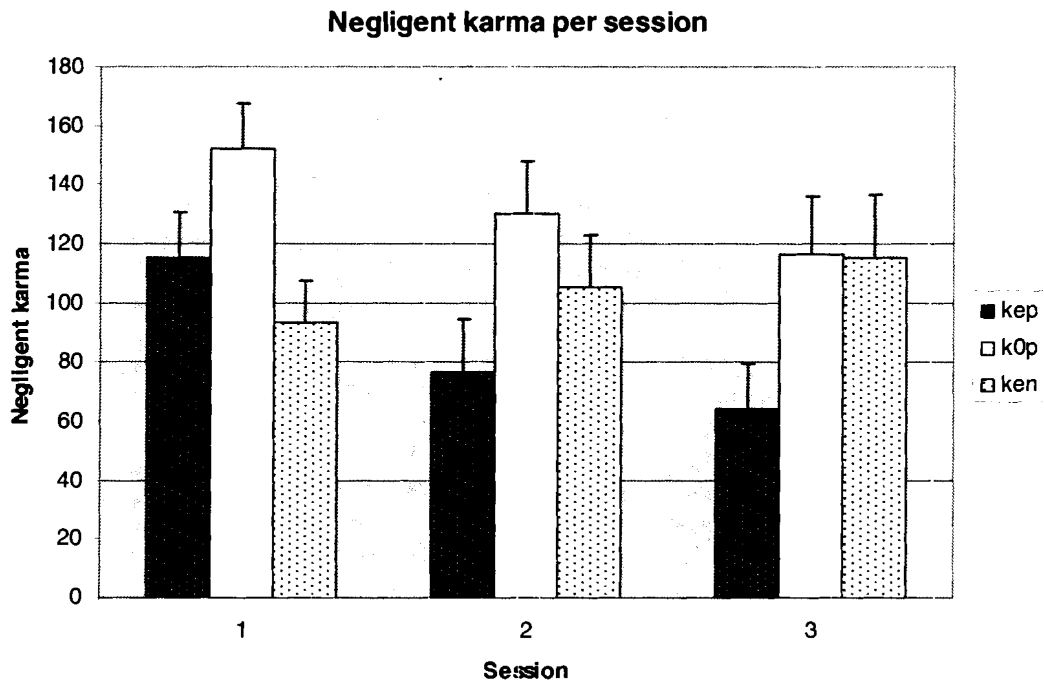

6.3.1 Analysis of aggressive karma

We used SPSS to run a one-way repeated measures ANOVA on the aggressive karma data and found the main effect (session) to be non-significant ($p = 0.692$), indicating that time alone did not have an effect on aggressive karma. Next, we found that the interaction between condition and session to be non-significant ($p = 0.166$), as summarized in Table 6.14. This indicates that no condition significantly changed their aggressive karma scores over time more so than others. Since there is no significant interaction or main effect, no further analysis of aggressive karma will be presented.

Table 6.11: Between-Subjects factors for aggressive karma data

		N
Condition	kOp	21
	ken	20
	kep	21

Table 6.12: Within-Subjects factors for aggressive karma data

Session	Dependent Variable
1	Session1
2	Session2
3	Session3

Table 6.13: Descriptive statistics for aggressive karma data

	Condition	Mean	Std. Deviation	N
Session1	kOp	197.2381	96.39186	21
	ken	255.7700	152.42112	20
	kep	133.8000	39.34819	21
	Total	194.6323	115.29123	62
Session2	kOp	211.1714	111.23998	21
	ken	237.4800	120.23481	20
	kep	160.8667	55.03883	21
	Total	202.6194	102.84431	62
Session3	kOp	210.9143	108.54996	21
	ken	224.8725	141.70463	20
	kep	154.3238	52.95611	21
	Total	196.2492	109.47621	62

Table 6.14: Tests of Within-Subjects effects on aggressive karma

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
session	Sphericity Assumed	2073.151	2	1036.576	.370	.692
session * Condition	Sphericity Assumed	18488.299	4	4622.075	1.650	.166
Error(session)	Sphericity Assumed	330583.355	118	2801.554		

Table 6.15: Tests of Within-Subjects contrasts for aggressive karma data

Source	session	Type III Sum of Squares	df	Mean Square	F	Sig.
session	Linear	37.547	1	37.547	.012	.914
	Quadratic	37.547	1	2035.604	.853	.359
session * Condition	Linear	15852.288	2	7926.144	2.464	.094
	Quadratic	2636.011	2	1318.006	.552	.578
Error(session)	Linear	189823.344	59	3217.430		
	Quadratic	140755.010	59	2385.678		

Measure: aggressive karma

Figure 6.12: Plot of aggressive karma over session

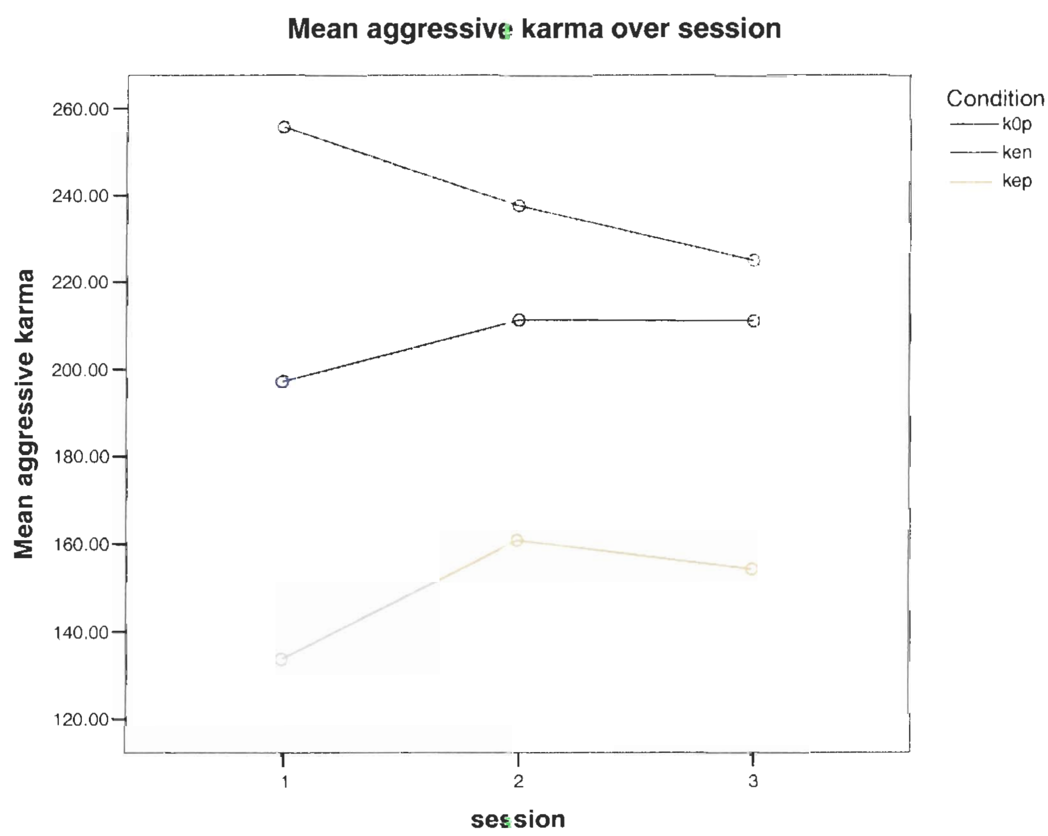

6.3.2 Analysis of negligent karma

Employing the same methods used in the analysis of mean *karma,* we found the main effect (session) to be significant (p = 0.037), indicating that session alone had an effect on negligent behaviour. Next, we found the interaction between condition and session also to be significant (p = 0.004), as summarized in Table 6.19. This indicates that at least one yet to be specified condition significantly changed their negligent karma scores over time more so than others. The expectation is for the *kep* group to significantly decrease their negligent behaviour over time, which is observable in Figure 6.11.

Table 6.16: Between-Subjects factors for negligent karma data

		N
Condition	kOp	21
	ken	20
	kep	21

Table 6.17: Within-Subjects factors for negligent karma data

Session	Dependent Variable
1	Session1
2	Session2
3	Session3

Table 6.18: Descriptive statistics for negligent karma data

	Condition	Mean	Std. Deviation	N
Session1	kOp	156.2667	77.06512	21
	ken	90.2300	62.62911	20
	kep	115.3619	70.54535	21
	Total	121.1097	74.49947	62
Session2	kOp	141.2571	80.99213	21
	ken	101.2000	83.52102	20
	kep	76.6381	82.83421	21
	Total	106.4484	85.47086	62
Session3	kOp	116.6976	87.75401	21
	ken	115.1375	97.28291	20
	kep	64.3714	68.54847	21
	Total	98.4710	87.28686	62

Table 6.19: Tests of Within-Subjects effects on negligent karma

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
session	Sphericity Assumed	15291.633	2	7645.817	3.382	.037
session * Condition	Sphericity Assumed	36393.101	4	9098.275	4.024	.004
Error(session)	Sphericity Assumed	266796.356	118	2260.986		

Table 6.20: Tests of Within-Subjects contrasts for negligent karma data

Source	session	Type III Sum of Squares	df	Mean Square	F	Sig.
session	Linear	14838.353	1	14838.353	4.410	.040
	Quadratic	453.280	1	453.280	.392	.534
session * Condition	Linear	34056.240	2	17028.120	5.061	.009
	Quadratic	2336.861	2	1168.431	1.009	.371
Error(session)	Linear	198507.616	59	3364.536		
	Quadratic	68288.740	59	1157.436		

Measure: negligent karma

Figure 6.13: Plot of mean negligent karma over session

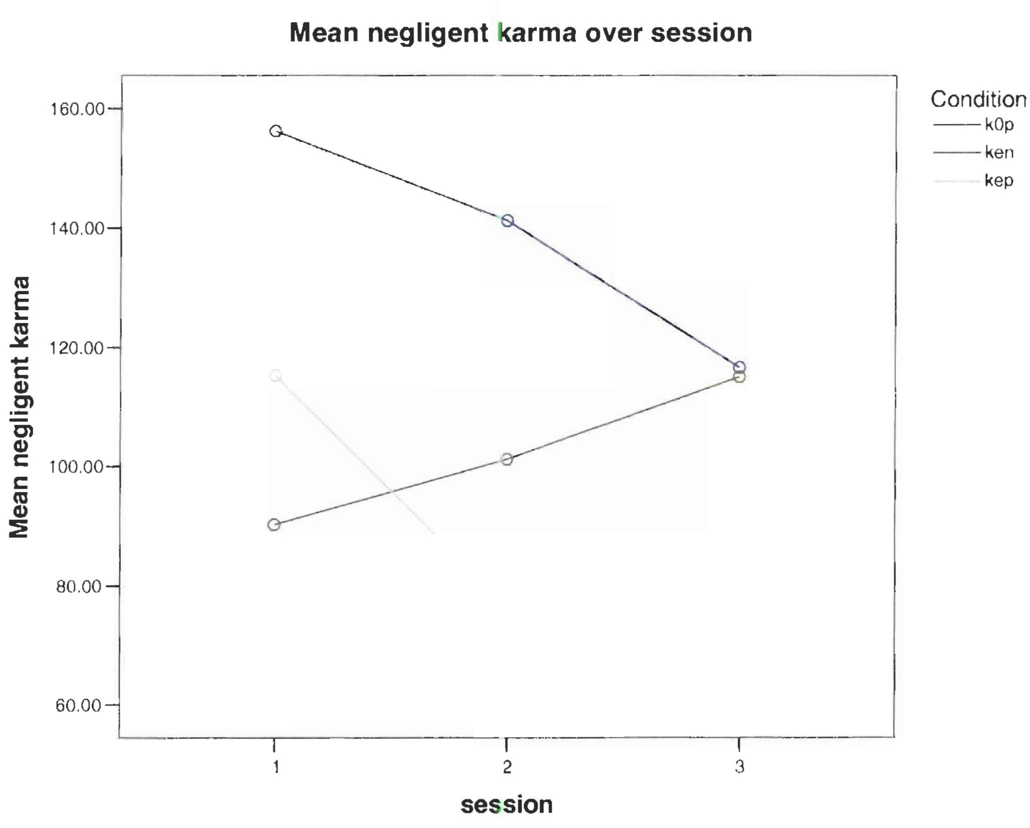

Given the significant session*condition interaction for negligent karma, we ran three separate one-way ANOVAs, one for each condition, to determine which group significantly changed their negligent behaviour over time. We found that the control and *kep* group significantly changed their negligent karma scores over time ($p = 0.037$ and $p < 0.0001$, respectively). This change was significantly linear for both the *kep* ($p = 0.001$) and *k0p* ($p = 0.031$) group. The summary of this analysis can be found in Table 6.21.

Table 6.21: Results of repeated measures ANOVA across conditions on negligent karma

Condition	Source	Type III Sum Of Squares	df	Mean Square	F	Sig.
kep	Session	29750.232	2	14875.116	12.060	.000
	Error	49338.194	40	1233.455		
k0p	Session	16759.159	2	8379.579	3.572	.037
	Error	93844.430	40	2346.111		
ken	Session	6233.189	2	3116.595	.958	.393
	Error	123613.733	38	3252.993		

Table 6.22: Tests of within-subjects contrasts for k0p negligent karma

Source	session	Type III Sum of Squares	df	Mean Square	F	Sig.
session	Linear	16439.950	1	16439.950	5.396	.031
	Quadratic	319.209	1	319.209	.194	.664
Error(session)	Linear	60931.401	20	3046.570		
	Quadratic	32913.028	20	1645.651		

Measure: k0p negligent karma

Table 6.23: Tests of within-subjects contrasts for kep negligent karma

Source	session	Type III Sum of Squares	df	Mean Square	F	Sig.
session	Linear	27300.301	1	27300.301	16.639	.001
	Quadratic	2449.931	1	2449.931	2.966	.100
Error(session)	Linear	32815.579	20	1640.779		
	Quadratic	16522.615	20	826.131		

Measure: kep negligent karma

Table 6.24: Scheffe comparisons of Session pairs for kOp negligent karma

(I) Session	(J) Session	Mean Difference (I-J)	Std. Error	Sig.	95% Confidence Interval	
					Lower Bound	Upper Bound
0	1	21.6348	24.36779	.676	-39.4350	82.7045
	2	35.3111	24.94123	.373	-27.1958	97.8180
1	2	13.6763	24.94123	.861	-48.8306	76.1832

Based on observed means, Dependent Variable: *kOp* negligent karma

Table 6.25: Scheffe comparisons of Session pairs for kep negligent karma

(I) Session	(J) Session	Mean Difference (I-J)	Std. Error	Sig.	95% Confidence Interval	
					Lower Bound	Upper Bound
0	1	38.7238	22.91256	.248	-18.7900	96.2376
	2	50.9905	22.91256	.093	-6.5234	108.5043
1	2	12.2667	22.91256	.867	-45.2472	69.7805

Based on observed means, Dependent Variable: *kep* negligent karma

The expectation was that the *kep* group's negligent karma would decrease over time more so than the control group. To determine whether the *kep* group changed their negligent karma over time more so than the control group, we ran another two-factor repeated measure ANOVA, but with the *ken* group removed from the negligent karma data. We were looking for a significant interaction effect between the *kep* and *kOp* groups, as we have already determined that the main effect is significant within each of these groups. What we found was that there was no significant interaction effect between the *kep* and control groups ($p = 0.442$). This means that the negligent karma rate of change is not significantly different between these groups.

Table 6.26: Tests of within-subjects effects on kep and k0p negligent karma

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
session	Sphericity Assumed	43555.634	2	21777.817	12.168	.000
session * Condition	Sphericity Assumed	2953.758	2	1476.879	.825	.442
Error(session)	Sphericity Assumed	143182.624	80	1789.783		

Of note, I ran a curve estimation analysis to determine the direction and magnitude of change within the *kep* and *k0p* groups. The trend analysis revealed that the *kep* group had the only set of significant coefficients, those being for the linear regression (Table 6.27). The *kep* group had a significant slope coefficient of -25.495 compared the control group's slope of -17.718. This indicates that even though both the *kep* and control group significantly decreased their negligent karma scores over time, the experimental group *kep* decreased, albeit insignificantly, their negligent behaviour at a greater rate.

Table 6.27: Linear coefficients for negligent karma scores across condition

		Unstandardized Coefficients		Standardized Coefficients		
Condition		B	Std. Error	Beta	t	Sig.
kep	Session	-25.495	11.404	-.275	-2.236	.029
	(Constant)	110.952	14.722		7.536	.000
k0p	Session	-17.718	12.373	-.175	-1.432	.157
	(Constant)	150.703	15.637		9.638	.000
ken	Session	10.807	12.712	.109	.850	.399
	(Constant)	93.734	16.224		5.777	.000

6.4 Analysis of the questionnaire

Each participant answered a questionnaire upon completion of the third session. Each question had a five-point Likert response scale, with Strongly Disagree being scored as 1 and Strongly Agree being scored as 5. A copy of the questionnaire can be found in Appendix B. Similar to the game-based data collection, there were occasions of data corruption in the questionnaire responses. For example, one participant was observed checking random boxes without reading the questions. This is in addition to not including those participants who tampered with the game. In total, we removed 9 questionnaires from the analysis. Furthermore, question 10 was not posed to the first pool of subjects because the need for this question was not determined until after the initial testing. This explains why question 10 has a lower response count than the other questions. To analyze the questionnaire, SPSS version 13 and the General Linear Model: Multivariate analysis was employed.

The analysis of the questionnaire employed the General Linear Model Multivariate procedure, with a single factor (condition) and 10 variables (questions). The GLM Multivariate procedure provides regression analysis and analysis of variance for multiple dependent variables by one or more factor variables or covariates. There were four main questions from the questionnaire that we were interested in:

- Aggressive play (for example, speeding up to deliver a big hit) is important to winning at Heads Up Hockey (Q1)
- Removing injured players from the lineup is not important to winning at Heads Up Hockey (Q3)

- The more I played Heads up Hockey, the less aggressive I played (Q8)
- If an injured player is not removed from the lineup, the rest of the team would start playing more poorly (for example, skate slower, shoot less hard) (Q9)

These four questions were intended to measure the level of awareness a participant had about the karmic teaching mechanism and in particular its emphasis on reducing aggressive and negligent behaviours. Questions 1 and 8 were on aggression and questions 3 and 9 were on negligence. The expectation was that the *kep* group would score lower on question 1 and higher on the other three as compared to the control group (Table 6.29). We expect the control group's responses for these four questions to be between the *kep* and *ken* group's scores.

Table 6.28: Descriptive statistics for the questionnaire results

Condition		Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10
kOp	Mean	4.19	3.38	3.00	3.19	4.95	3.67	4.19	1.81	3.05	4.75
	N	21	21	21	21	21	21	21	21	21	16
	Std. Dev.	.750	1.284	1.789	1.537	.218	1.278	.750	1.250	1.396	.577
ken	Mean	3.55	3.75	2.20	2.65	4.70	2.70	3.80	2.35	2.90	4.19
	N	20	20	20	20	20	20	20	20	20	16
	Std. Dev.	1.276	1.209	1.399	1.348	.923	1.720	1.322	1.387	1.165	1.377
kep	Mean	3.39	3.13	2.38	2.42	4.58	3.21	4.17	2.67	3.83	3.74
	N	23	24	24	24	24	24	24	24	24	19
	Std. Dev.	1.234	1.424	1.715	1.100	1.018	1.474	.761	1.341	.963	1.558
Total	Mean	3.70	3.40	2.52	2.74	4.74	3.20	4.06	2.29	3.29	4.20
	N	64	65	65	65	65	65	65	65	65	51
	Std. Dev.	1.150	1.321	1.659	1.350	.815	1.523	.966	1.355	1.234	1.312

Based in a Likert-scale from 1 (Strongly disagree) to 5 (Strongly agree)

Table 6.29: Expectation of mean score on aggressive and negligent questions

Question	kep	ken
Q1	< kOp	> kOp
Q3	< kOp	> kOp
Q8	> kOp	< kOp
Q9	> kOp	< kOp

The results show that the *kep* group responded as expected for all four questions. However, the *ken* group only responded as expected for question 9. This indicates that the *kep* group had some knowledge that the game was trying to change their aggressive and negligent behaviours. On the other hand, the *ken* group did not respond as expected, which is also evident in the analysis of the *ken* group's composite karma scores. The next step is to determine if any of these differences are significant.

Table 6.30: Average scores on aggressive and negligent questions

Question	kep	kOp	ken
Q1	3.39	4.19	3.55
Q3	2.38	3.00	2.20
Q8	2.67	1.81	2.35
Q9	3.83	3.05	2.90

Based in a Likert-scale from 1 (Strongly disagree) to 5 (Strongly agree)

The GLM Multivariate analysis produced a pair-wise comparison of each condition within each question. The results reveal that only one pair has a p value less than 0.05, and that is in question 8 between the *kep* and *kOp* group ($p = 0.030$). Table 6.31 summarizes the results of the pair-wise comparison for the aggressive and negligent questions.

Table 6.31: Subset of the pairwise comparison table from the questionnaire analysis

Dependent Variable	(I) Condition	(J) Condition	Std. Error	Sig.
Q1	kOp	ken	.400	.166
		kep	.389	.086
	kep	ken	.400	.166
Q3	kOp	ken	.579	.201
		kep	.562	.448
	kep	ken	.579	.201
Q8	kOp	ken	.447	.332
		kep	.434	.030
	kep	ken	.447	.332
Q9	kOp	ken	.392	.635
		kep	.381	.067
	kep	ken	.381	.022

Full table can be found in Appendix D

The lack of significant results for this subset of questions may be related to several reports from participants (some as young as 10 years old) that the questions were too difficult to understand. In addition, through participant observation, two recorded post-study interviews and off the record conversations, it became clear that players were mostly unaware of changes in their gameplay behaviour. This would have an effect on the questionnaire because a majority of the questions required players to reflect upon and identify their behaviour during gameplay. However, there were participants who were able to understand what the game was trying to teach, exemplified by one *kep* participant (S0) who improved his *karma* scores after each session.

> **Q**: Can you tell me what you thought was the winning strategy at Heads Up Hockey?
>
> **S0**: You have to forget about trying to hit the opponent and just skate really fast around, and forget about trying to do fancy stuff. Just keep peppering the goalie with shots
>
> . . .

> **Q**: Do you think Heads Up Hockey is trying to teach you anything?
>
> **S0**: Yeah, probably about the brain injuries. That concussion can affect your play and everything that you do.

Even though the karmic mechanism was aimed at changing aggressive and negligent behaviours, the context of the educational scenario was concussion in hockey. Participant S0 was able to understand that trying to hit the opposition at every opportunity was not the proper strategy. He was also able to ascertain that the game was concerned with concussive injury in hockey. However, participant S0 was more of an exception than the rule as evidenced by the outcome of the questionnaire. Another *kep* participant (S1) also improved his *karma* scores after each session, but was not able to grasp what *Heads Up Hockey* was trying to accomplish.

> **Q**: If you were to pick one thing, what was the most important thing to winning at Heads Up Hockey?
>
> **S1**: Having a good lineup
>
> ...
>
> **Q**: Did you get the feeling that the game was trying to teach anything?
>
> **S1**: I'd think it was trying to teach me strategy and how to use players well and not over work them too much

Participant S1 came close to understanding and recognizing his own gameplay behaviour changes, especially as it relates to identifying negligent behaviour. However, there was a perceptible gulf between the actual behaviour change and his ability to correctly identifying the change. This disconnect is most evident within the *kep* group since they ideally would be able to consciously connect aggressive and negligent play with losing. The implications of the participants' lack of awareness regarding their own behaviour will be discussed in the coming chapter. In addition, I will present some changes to the design that could improve the efficiency of the *Heads Up Hockey* karmic teaching mechanism.

CHAPTER 7: CONCLUSION

This thesis describes the design, implementation and assessment of the educational sports-action video game *Heads Up Hockey*. The context of this game is concussive injury in hockey, with the intent to reduce the behaviours that can lead to concussions. The question I asked was how educational digital games can be developed in a way that preserves inherently enjoyable gameplay aspects at the same time as reliably changing learner outcomes. After examining past educational gaming efforts, the answer reached in this thesis was to use an implicit teaching mechanism that situates the learner in a role where they can learn through authentic challenges. In order to test this concept and advance the methods of concussion education, we designed and built *Heads Up Hockey*, with an implicit teaching mechanism that we labelled 'karma'.

The results of the assessment study clearly showed that the implicit teaching mechanism performed as intended. The composite behaviour score *karma* significantly improved over time amongst the experimental participants as compared to the control group. Reduction in negligent behaviour was the primary reason for the overall improvement in *karma* with the experimental group. The individual negative components of the composite karma score (aggressive and negligent) did not significantly decrease, as they were designed to work in concert with each other and positive karma to affect behaviour. This is why the primary evaluation of the effectiveness of the karmic teaching mechanism is based on the composite karma score.

7.1 Discussion of results

The study and analysis aimed to answer three main questions: a) did the *kep* group improve their composite karma score over the control group b) did the *kep* group reduce their negligent karma totals over the control group and c) did the *kep* group reduce their aggressive karma totals over the control group. We found significant evidence that the *kep* group improved their *karma* scores versus the control group, in addition to significantly reducing their negligent behaviour. However, there was no significant interaction effect between the *kep* and control group for the negligent karma data, indicated that these groups were not significantly different in their rates of negligent karma change. In addition, no significant interaction or main effect was found in the aggressive karma data, which means there was no significant change in aggressive karma among the three groups (*kep*, *ken* and *k0p*). As previously mentioned, aggressive and negligent karma were designed to work with each other to affect behaviour change, thus separating them from the *karma* score ends up diluting their effect. The analysis of aggressive and negligent karma sheds light on the proportion of influence each had on the significant change in experimental composite karma scores.

The *kep* group responded to the questionnaire in the expected manner as compared to the control group, however only one result was significant (question 8). Other than methodological issues, one explanation for the lack of more significant questionnaire results may be the fact that the implicit nature of the instruction did not provide enough scaffolding for the learner to be able to properly reflect and identify their changed behaviour. However, since there was no external scaffolding or facilitation, we were able to rule out those as exogenous factors and more clearly evaluate the implicit teaching mechanism, which was the primary goal of the study.

In terms of the overall karmic mechanism, *Heads Up Hockey* was able to positively change the behaviour of the *kep* group. This is an interesting result because there was no explicit communication of the karma system to the participants. The reason why we masked the teaching mechanism is to emphasize fun and playability, which would increase player engagement and in turn provide greater opportunity for learning. Keeping learners engaged with a new technology learning medium is an important first step. The next step is to have an efficient mechanism for transferring content, skills or behaviours. The version of *Heads Up Hockey* tested for in this project was built to transfer behaviours. In particular, the game (under the *kep* condition) attempted to teach players that aggressive and negligent behaviour would not lead to winning. We were successful in reducing these behaviours, even though participants were unable to fully grasp and recognize the change. However, because we have a record of the behaviour change, we have the tools to help future participants reflect upon and learn from the experience.

To accomplish our educational goals, we connected the avoidance of aggressive behaviour with an increased chance of being competitive, which is a primary motivation for playing games (Deci & Ryan, 1985). However, there is a risk in this approach for non-competitive participants, as they are more likely to disengage from the game as a result of losing and thus not be exposed to the karmic teaching mechanism. This phenomenon of disengagement occurred several times, most notably under the more difficult *kep* condition. Nonetheless, the results show that the designed connection between behaviour and winning, via karma, achieved the desired behavioural outcome. The downside is that there remains a gap within the participant's mind as to why they started winning or why they changed their behaviour. Some thoughts on how to bridge that gap will be presented in the next section.

There is an open question about whether participants should be told that their behaviour change was motivated by brain injury concerns. The danger is that if we emphasize the connection between aggressive play and concussion then they might not internalize the behaviour change as readily. Aggressive play is often coached as one of the tenets of a successful hockey team and is how a player usually wins at other hockey video games. Anything that contradicts that philosophy needs to be framed in equally strong terms (e.g. if you play aggressively you will lose.) It should be noted that our goal is not to remove all aggressive play from hockey. The goal is to reduce unsafe aggressive play and in turn reduce the rate of concussive injury by increasing players' awareness of the consequences of playing with concussive injuries. In addition, we are trying to instil the mindset that aggressive play will not only hurt your own performance (e.g. receiving penalties), but the team's as well (e.g. the team gets scored on during the ensuing power play). The best way to accomplish these goals within the simulation game context is to present aggressive and negative behaviour as contributors to losing, since losing is an outcome all competitive hockey players want to avoid.

The version of version of *Heads Up Hockey* tested here was not designed to teach about concussion. The version of the game tested for this project was aimed at curbing the behaviour that could lead to concussive injuries. Nonetheless, awareness of concussion was raised due to the occurrence of the injury during gameplay and its role in affecting game outcomes. Some of the questionnaire and interview responses indicated that participants were aware of the brain injury context within the game. While it was unlikely that the participants learned anything specific about concussions, their awareness level was raised. This is a positive first step towards being able to more effectively educate youth hockey players about the dangers of concussion, especially considering the main thrust of the education was not concerned with explicit concussion knowledge.

The purpose of designing and implementing an implicit teaching mechanism (*karma*) was to keep the engagement of the game as high as possible. Another reason was that by forcing a participant to struggle through the game and discover the proper behaviours on their own terms, we hoped to increase the durability and recall of the new game playing behaviours. Unfortunately we have no way of reporting on behaviour recall from our group of participants. We are able to report that the game was highly enjoyable for the players, which shows that our design approach of focusing on gameplay was successful. The original design could have made *karma* more explicit (e.g. include a *karma* meter), which may have resulted in the participants acquiring the new behaviours more quickly. The potential problem with being explicit is that a participant would not have to work as hard to figure out the key to the game. The player could circumvent the problem solving process that may eventually yield the proper behaviours needed to win. By intrinsically motivating a player to solve the karmic mechanism on their own terms, their solution becomes more personally meaningful and retained for a longer period (Perkins, 1991). In addition, a *karma* meter would immediately give away that the game is trying to reduce certain in-game behaviours, such as hitting. By being that apparent, we run the risk of alienating those players who find it 'boring' having to play it safe.

7.2 Possible alterations to the study

One of the problems of the study was that the *ken* group likely experienced a floor effect. It was predicted that the *ken* participants would continually decrease their *karma* scores over time, which we did not see. One alteration to the study that could alleviate these floor effect concerns would be to improve the normalization of the *karma* distributions. To accomplish this, we could change the weight of each karma reward and penalty depending on condition. The problem with this approach is that then we lose symmetry of gameplay among the conditions. That is, an aggressive hit would mean different things karma-wise under different conditions. A better solution would be to construct a secondary behaviour scale based on *karma* that does not have a floor or ceiling (i.e. an unlimited range). The actual effects of *karma* still need to be capped at 100 and -100 because anything beyond those limits would impair gameplay by assigning unrealistic abilities to players. This new behaviour scale would eliminate the floor and ceiling effects in addition to capturing a more normalized distribution of *karma* scores. In turn, we would be able to better compare the differences between each group's behaviour and hopefully see a significant difference between the *ken* and *k0p* groups.

Another issue is the linear relationship (Figure 6.1 and Figure 6.2) between *karma* and goal differential. The main driver of this correlation is the linear relationship between *karma* and player ability (e.g. shotpower, accuracy, etc.) As *karma* increased, ability linearly increased and vice versa. For the purposes of this study, this design proved effective. However, there were still some participants who could not figure out the game, or gave up after several losses. The loss of these participants was due in part to the difficulty of recovering from a poor *karma* score. Therefore, if we changed the relationship between *karma* and ability to be cubic, players would be able to more quickly recover from extreme *karma* scores. More specifically, as *karma* increased through the positive range, it would become more difficult to increase ability. In addition, accumulations of negative *karma*, if left unscaled, would have a greater effect. If we applied the same strategy across the negative range (but inversed), we would see a greater tendency for *karma* scores to hover around zero. This would allow players to recover from the extreme negative *karma* ranges. This change should be done in conjunction with a separate behaviour scale, so we only end up scaling *karma* scores and not behaviours (i.e. an aggressive hit should still be an aggressive hit regardless of *karma* score).

Figure 7.1: Possible relationships between *karma* and avatar ability

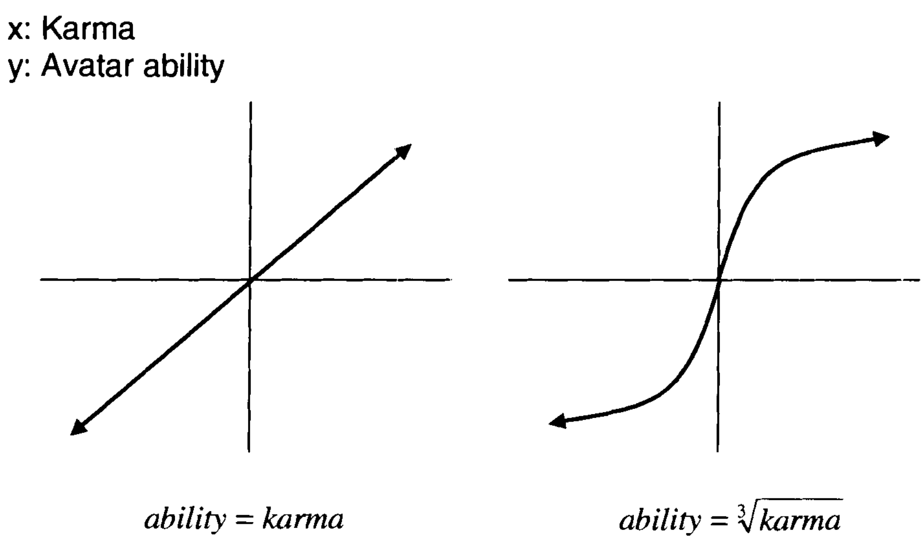

Further improvements to the learning tool need to take into account the magnitude of the *karma* penalties to ensure they corroborate with the learning goals. For example, we need to ensure that the *karma* penalty of charging is strong enough to deter that particular act over the long term. In addition, the differences in magnitudes between aggressive acts need to be considered, especially if the simple sum of aggressive and negligent karma penalties is to be used to evaluate behaviour change. This brings up another issue of counting individual acts of aggression and negligence. Currently, there is no simple way to examine the occurrence rates of a specific aggressive or negligent act. Separate counters for each aggressive and negligent act should be established, and this could be used in place of simple sums of karma penalties for evaluating behaviour change.

A significant reduction of aggressive behaviour alone was not found with the existing setup of *Heads Up Hockey*. Reducing aggressive acts will be the most important aspect of future versions of the game, as negligent behaviours are more within the domain of trainers and coaches. Therefore, future designs of the game need to change the negligent karma system to be an influencer on aggression. Concussions during gameplay are still important to changing the behaviour of players as they are the ultimate consequence of aggressive behaviour, even though it is usually the other team that would incur a concussion as a result of a participant's aggressive play. The number of aggressive acts may need to be expanded to include boarding hits and attempted charging hits. The penalties for each aggressive act may also need to be increased, depending on how the other changes to the karma system affect aggressive karma.

Without a negligent karma system, there would be no way to penalize leaving injured players in the lineup. Even though the focus is now on reducing aggressive acts we should still give some attention to negligent acts. One possible solution is to continue to record negligent karma, but have its effect be an increase in the magnitude of aggressive karma penalties. The downside of this solution is that if a player is able to avoid all aggressive acts, then there will be no noticeable difference in gameplay (except for the concussion indicator and disrupted player control). However, since our focus may no longer be on negligent acts, this solution should be acceptable for a program that focuses on reducing aggressive behaviour.

With respect to the questionnaire, one reason why there was a lack of significant results was because no participant was explicitly given any indication as to whether their behaviours changed. When they were forced to reflect upon their experience for the first time, there was not enough time or dialogue to comprehend what they just experienced (the test was administered seconds after the completion of their fifteenth game). Therefore, the level of comprehension needed to show clear differences between groups was not possible. A five minute discussion amongst the experimental participants on the questionnaire's topics would allow for both group and individual reflection. This kind of simple exchange would likely be sufficient to bring out some of the differences between groups. More complex discussion scaffolding could be used, such as a full debriefing by a facilitator that included individual behaviour scores. There is a potential problem that a minority of participants may dominate the discussion, but it is the facilitator's job to intervene and elicit other perspectives.

In terms of the validity of the questionnaire, there was no pre-testing of each group to ensure normal distribution. The design of the study made it possible to not have to administer a pre-test because of the inclusion of a control group. The experimental group's results could be compared to the control group to determine the effects of the intervention. However, this design relies on random assignment and normal distribution of the groups. With sample sizes of 20 there is a possibility that the participants were not normally distributed. A preexperiment baseline test could have been used to verify the normalcy of the groups and provide direction on how to adjust the results, if necessary.

Another improvement to the study would be to increase the overall consistency of the testing experience for each participant. There was little consistency in environment, number of co-participants and computing platforms across subject pools. That said, despite the lack of stringent consistency for all *Heads Up Hockey* participants, the differences are not as noticeable as they might seem due to the engaging properties of the game. When players engage with a game, they enter a "magic circle" (Salen & Zimmerman, 2004) that shuts out a lot of the surrounding world. Therefore, the variable environments likely did not have a significant effect on those participants who were engaged with the game. With respect to the computing speed disparity, I do not believe these differences would have a significant effect on the results. For one, the kind of behaviours we are looking for are not a function of the overall speed of the game, especially the negligent behaviours. Also, the actual speed difference between our two computer platforms (laptops and desktops) is small enough that the actual game experience is perceptually identical.

7.3 Future directions

The result of this project opens up a wide field of possibilities. We have shown that a player can exhibit new behaviours through the play of an engaging video game without explicitly making known the teaching mechanism. Our particular approach was concerned with health related behaviours, of which there are many possible reapplications of this approach (e.g. smoking, diabetes self-care, diet, and so on). However, to reapply this approach to other behaviours there needs to be two major considerations. First, the game itself needs to be as engaging and fun as possible. Engaging media, such as *Heads Up Hockey*, has been argued to have the potential to be more effective than traditional media (e.g. the lecture) at transferring knowledge (Foreman, 2003). Second, the participants need to be subjected to some kind of feedback to ensure that they understand how and why their game playing behaviour changed. The results from this study show that the game needed a facilitating influence or a modified mechanism of reward and punishment to ensure that the participants fully understood what they were learning. If these issues are not addressed, the likelihood that new behaviours will transfer over into real life scenarios is lessened.

Looking again at Egenfeldt-Nielson's generational perspective on educational video games, we find that *Heads Up Hockey* fits in on the "smart" end of second generation educational games. This means that the game is effective at connecting the game actions with the overall educational aims. This connection between gameplay and education is found within a host of other successful educational titles (e.g. *Carmen San Diego* and *Oregon Trail*). It does not reach third generational status because of the lack of a community perspective. The community perspective is needed to cement the behaviours that are being learned by participants. In this study, the game was the primary source of education. Interaction between peers was minimal, and the experimenters were instructed to say as little as possible to the players. As a result, the game was able to affect behaviours, but the participants were unable to grasp the change. By introducing a third generation perspective to the study, I would expect to see more self-recognition of the behaviour changes, further comprehension as to why these changes occurred, and higher awareness about the dangers of aggressive and negligent play in hockey. In turn, these changes increase the likelihood of knowledge transfer from the video game context to actual ice hockey situations.

One way to achieve a third generational game is to introduce a facilitator or tools for reflection. The job of the facilitator would be to debrief players on their play, describe what the game was about, and show them why they won or lost. This opportunity could also be used to impart concussion knowledge or other lessons. In addition, a third generational game should encourage interaction between game players. The best way to do that would be to have a multiplayer version of the game. The facilitator could also encourage discussion outside of the game playing. Online multiplayer capability, Internet sites, web chats and blogs are all tools currently used to produce a gaming community and should be considered.

Keeping in mind that a major goal of this project is to address concussion in hockey, future efforts should be directed towards testing *Heads Up Hockey* on youth hockey players. The study should attempt to capture differences in actual on ice aggressive and negligent behaviours but with a special focus on aggressive behaviours. Players are not in control of who remains in the lineup, therefore negligent behaviours are usually the domain of a coach or trainer.

The lessons learned from this project can be implemented in a new version of *Heads Up Hockey*, resulting in a more effective learning experience. In addition, a new study on youth hockey players should look at including a third party facilitator, or ensuring there is some sort of community discussion around the gameplay. The ThinkFirst Canada evaluation study (Cook *et al.*, 2003) provides a good starting point on how to structure the experimental design. Retention was important to that study, and should be included in an updated *Heads Up Hockey* experiment.

Another challenge for a future experiment using *Heads Up Hockey* is how to separate and isolate the effects of negligent and aggressive karma. A further understanding of how participants respond to these components of *karma* will be important to the design of new teaching mechanism. Towards this goal, *Heads Up Hockey* records a play-by-play file of events that occur in the game (Appendix A). To help understand what the player's thought processes were, these play-by-play files could be used in conjunction with recorded video of the gameplay using a think-aloud protocol (Nielson, 1993). The goal would be to see the level to which the participants recognize the various effects of aggressive and negligent karma. This information would then be fed into the redesign process of the teaching mechanism, thereby fostering an iterative evolution of our karma system.

Finally, looking at the wider scope of educational gaming, it is clear that the field remains limited on several levels (Egenfeldt-Nielsen, 2005). The key to breathing some life into this genre resides with the level of engagement future educational games exhibit. In addition, the teaching mechanism needs to be tied to that engagement in order to deliver powerful educational gaming experiences. This project provides one successful blueprint for a specific application (health behaviours) within a particular game genre (sports/action). While our approach may not extend beyond these genres, there are enough unique combinations between health behaviour and action games for there to be another application. In terms of design principles, *Heads Up Hockey* and other games that follow in its mould should rely on fun-first game-play melding cleanly with the teaching mechanism. A player should learn through playing rather than having to learn in order to get to gameplay. The two, playing and learning, should be intertwined.

APPENDICES

Appendix A: One period of play-by-play data

#eid:	Per	Time	Event	Team	Description
1	1	2:00	FACEOFF	N/A	Away wins
2	1	1:54	SHOT	Away	Away shoots
3	1	1:48	SHOT	Home	Home shoots
4	1	1:35	SHOT	Away	Away shoots
5	1	1:26	SHOT	Home	Home shoots
6	1	1:21	FACEOFF	N/A	Away wins
7	1	1:10	SHOT	Away	Away shoots
8	1	1:04	SHOT	Home	Home shoots
9	1	1:03	KARMA	Home	Start: 29 Delta: -11 Final: 18
10	1	1:03	HITTARG	Home	Goalie
11	1	0:51	SHOT	Away	Away shoots
12	1	0:49	INJURY	Home	Home team has sustained an injury
13	1	0:45	SHOT	Home	Home shoots
14	1	0:41	LINECHG	Home	Home is changing lines
15	1	0:35	SHOT	Away	Away shoots
16	1	0:33	LINECHG	Home	Home is changing lines
17	1	0:33	LINECHG	Home	Home is changing lines
18	1	0:32	LINECHG	Home	Home is changing lines
19	1	0:31	SHOT	Away	Away shoots
20	1	0:21	SHOT	Home	Home shoots
21	1	0:20	SHOT	Home	Home shoots
22	1	0:16	HITFREQ	Home	First hit in a row
23	1	0:16	KARMA	Home	Start: 9 Delta: -7 Final: 2
24	1	0:16	HITTARG	Home	Open player
25	1	0:15	FACEOFF	N/A	Away wins
26	1	0:07	SHOT	Away	Away shoots

Appendix B: Post-play questionnaire

Five possible responses for each question, ranging from Strongly Agree to Strongly Disagree

Name: ____ **Age**: ________

Gender: Male or Female (Circle one)

For each of the statements below, please indicate the extent of your agreement or disagreement by placing a tick in the appropriate column.

Please answer all questions based on your experience of playing Heads Up Hockey.

1. Aggressive play (for example, speeding up to deliver a big hit) is important to winning at *Heads Up Hockey*

2. Changing lines every 25 seconds or so is important to winning at *Heads Up Hockey*

3. Removing injured players from the lineup is not important to winning at *Heads Up Hockey*

4. Chippy play (for example, hitting a player that doesn't have the puck) improves my chances of winning at *Heads Up Hockey*

5. Getting used to the game controls is important to winning

6. The more I played *Heads Up Hockey*, the more often I removed injured players from the lineup

7. The more tired a player was in *Heads Up Hockey*, the more likely he was to get injured

8. The more I played *Heads Up Hockey*, the less aggressive I played

9. If an injured player is not removed from the lineup, the rest of the team would start playing more poorly (for example, skate slower, shoot less hard)

10. I tried my absolute best to win at *Heads Up Hockey*.

Appendix C: Data subset

Subset of data that Heads Up Hockey records for each participant

Stat	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10	G11	G12	G13	G14	G15
Home Goals: 1st	0	1	1	0	1	0	0	3	1	1	1	0	3	2	0
Home Goals: 2nd	1	0	1	0	0	0	0	2	0	2	2	1	2	2	0
Home Goals: 3rd	0	0	0	0	0	0	0	1	1	0	1	0	3	2	1
Away Goals: 1st	3	2	1	3	2	1	1	0	0	3	3	2	1	1	0
Away Goals: 2nd	2	1	1	4	6	1	2	1	3	3	1	1	1	1	2
Away Goals: 3rd	5	5	1	2	1	4	4	1	0	0	1	2	1	1	2
Home Shots: 1st	4	3	6	2	1	4	5	7	6	10	9	6	10	11	4
Home Shots: 2nd	6	5	2	1	4	4	5	9	7	7	10	5	12	10	12
Home Shots: 3rd	3	1	2	2	0	5	3	8	4	2	6	2	11	9	8
Away Shots: 1st	12	8	9	10	16	9	10	13	4	12	11	12	7	8	7
Away Shots: 2nd	10	11	10	12	12	9	11	14	10	6	6	11	3	9	7
Away Shots: 3rd	11	20	12	14	8	11	12	10	14	9	11	10	8	8	14
Home attack time	91	57	102	86	92	123	100	101	95	98	107	103	99	111	113
Home passes att.	28	37	28	34	34	32	40	30	27	34	29	34	28	36	36
Home passes comp.	5	13	10	10	10	18	13	12	9	13	7	10	13	12	18
Away faceoffs won	12	10	8	9	7	11	5	5	7	7	3	4	6	4	6
Away attack time	184	223	177	211	198	158	191	183	188	196	163	172	168	163	174
Home shots	13	9	10	5	5	13	13	24	17	19	25	13	33	30	24
Away shots	33	39	31	36	36	29	33	37	28	27	28	33	23	25	28
Puck carrier: Total	4	4	6	2	0	6	3	1	4	5	1	2	7	8	1
Interference:Total	7	3	2	6	8	2	7	8	2	6	8	7	8	3	4
Goalie:Total	5	1	4	4	2	8	7	3	5	8	6	5	2	4	2
Firsts:Total	10	7	6	5	8	8	10	9	5	10	9	8	14	11	3
Seconds:Total	0	0	2	3	0	0	0	0	1	1	0	1	1	0	1
Thirds:Total	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
Excess:Total	7	5	6	7	5	5	4	2	2	2	3	5	6	6	2
#Inj:Total	3	5	1	4	4	5	3	3	1	2	1	3	2	2	2
Avg. Karma:1st	-13	-23	13	-71	-55	-64	-9	15	8	-1	14	-6	27	4	-27
Avg. Karma:2nd	-49	-91	-9	-69	-92	-88	-20	49	-22	7	26	-9	3	46	-50
Avg. Karma:3rd	-72	-90	-50	-95	-63	-87	-19	51	14	3	15	-55	-25	52	-8
Avg. Karma:Total	-44	-68	-15	-78	-70	-79	-16	38	0	3	18	-23	1	34	-28
Avg. Energy 1:Total	57	67	68	46	47	60	77	89	89	86	92	86	92	93	81
Avg. Energy 2:Total	91	83	88	85	83	77	83	91	93	91	89	89	91	90	86
Avg. Energy 3:Total	96	79	97	88	87	77	92	95	90	90	93	88	92	95	92
Negligent Karma	-35	-206	-13	-116	-74	-76	-10	-12	-8	-10	-4	-37	-6	0	-20
Aggressive Karma	-181	-87	-197	-265	-133	-157	-170	-111	-124	-174	-155	-199	-177	-131	-148

Appendix D: Questionnaire analysis

Pairwise comparison of all groups within each question

Dependent Variable	(I) Condition	(J) Condition	Mean Difference (I-J)	Std. Error	Sig.	95% Confidence Interval for Difference	
						Lower Bound	Upper Bound
Q1	KOP	KEN	.563	.400	.166	-.242	1.367
		KEP	.681	.389	.086	-.101	1.462
	KEN	KOP	-.563	.400	.166	-1.367	.242
		KEP	.118	.389	.763	-.664	.900
Q2	KOP	KEN	-.125	.475	.793	-1.080	.830
		KEP	.556	.461	.234	-.372	1.483
	KEN	KOP	.125	.475	.793	-.830	1.080
		KEP	.681	.461	.147	-.247	1.608
Q3	KOP	KEN	.750	.579	.201	-.414	1.914
		KEP	.431	.562	.448	-.701	1.562
	KEN	KOP	-.750	.579	.201	-1.914	.414
		KEP	-.319	.562	.573	-1.451	.812
Q4	KOP	KEN	.813	.463	.086	-.119	1.744
		KEP	.688	.450	.133	-.218	1.593
	KEN	KOP	-.813	.463	.086	-1.744	.119
		KEP	-.125	.450	.782	-1.031	.781
Q5	KOP	KEN	.250	.321	.440	-.396	.896
		KEP	.493	.312	.121	-.135	1.121
	KEN	KOP	-.250	.321	.440	-.896	.396
		KEP	.243	.312	.440	-.385	.871
Q6	KOP	KEN	.938	.518	.077	-.105	1.980
		KEP	.278	.504	.584	-.736	1.291
	KEN	KOP	-.938	.518	.077	-1.980	.105
		KEP	-.660	.504	.197	-1.673	.354
Q7	KOP	KEN	.438	.335	.198	-.237	1.112
		KEP	-.215	.326	.512	-.871	.440
	KEN	KOP	-.438	.335	.198	-1.112	.237
		KEP	-.653	.326	.051	-1.309	.003
Q8	KOP	KEN	-.438	.447	.332	-1.336	.461
		KEP	-.972(*)	.434	.030	-1.846	-.099
	KEN	KOP	.438	.447	.332	-.461	1.336
		KEP	-.535	.434	.224	-1.408	.339
Q9	KOP	KEN	.188	.392	.635	-.601	.976
		KEP	-.715	.381	.067	-1.482	.051
	KEN	KOP	-.188	.392	.635	-.976	.601
		KEP	-.903(*)	.381	.022	-1.669	-.136
Q10	KOP	KEN	.563	.429	.196	-.301	1.426
		KEP	.861(*)	.417	.044	.022	1.700
	KEN	KOP	-.563	.429	.196	-1.426	.301
		KEP	.299	.417	.477	-.540	1.137

Appendix E: Description of player attributes

Attribute	Type	Description
Name	String	First and last name of player for use in UI
Mass	Integer	Used in speed and collision equations
Height	String	For display purposes only (recorded in feet and inches)
Skating speed	Integer	Determines how fast a player can move around the ice
Shot accuracy	Integer	Effects whether a shot hits the desired target or not
Faceoff	Integer	Controls how quick players are at winning faceoffs
Acceleration	Integer	Determines how quickly a player can turn
Agility	Integer	Controls how quickly a player can turn
Shot power	Integer	Effects the shot velocity
Aggressiveness	Integer	Used in the calculation of injury likelihood
Jersey number	String	For display purposes only
Energy	Integer	Effects all other attributes, and is decreased through play and increased through rest
Brain injury level	Integer	The current concussive state of a player
Injury resistance	Integer	Is used in the likelihood of injury calculation
Team	String	Can be either Home or Away

REFERENCE LIST

Adams, P. (1998). Teaching and learning with simcity 2000. *Journal of Geography, 97*(2), 47-55.

Asgari, M., & Kaufman, D. (2004, June 14-17). *Relationships among computer games, fantasy, and learning.* Paper presented at the Educating Imaginative Minds: 2nd Annual Conference on Imagination and Education, Vancouver, B.C.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. (1997). Deciding advantageously before knowing the advantageous strategy. *Science, 275*, 1293-1295.

Bertsche, D., Crawford, C., & Macadam, S. E. (1996). Is simulation better than experience? *McKinsey Quarterly*(1), 50-57.

Blossom, J., & Michaud, C. (1999). Postmortem: Lucaslearning's star wars droidworks. Retrieved October 23, 2005, from http://www.gamasutra.com/features/19990813/droidworks_01.htm

Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher, 18*(1), 32-42.

Brown, S. J., Lieberman, D. A., Gemeny, B. A., Fan, Y. C., Wilson, D. M., & Pasta, D. J. (1997). Educational video game for juvenile diabetes: Results of a controlled trial. *Medical Informatics, 22*(1), 77-89.

Buckingham, D., & Scanlon, M. (2003). *Education, entertainment and learning in the home.* Buckingham: Open University Press.

Burton, J., Moore, D., & Magliaro, S. (2004). Behaviorism and instructional technology. In D. Jonassen (Ed.), *Handbook of research on educational communications and technology* (2nd ed., pp. 3-36). Mahwah, NJ: Lawrence Erlbaum.

Ciavarro, C., Meanley, J., Bizzocchi, J., & Goodman, D. (2005). *Embedding educational content between gameplay: An example from a sports action videogame.* Paper presented at the ED-MEDIA 2005, Montreal, Que.

Cockburn, A., & McKenzie, B. (2002). Evaluating the effectiveness of spatial memory in 2d and 3d physical and virutal environments, *CHI'2002 Conference on Human Factors in Computing Systems.* Minneapolis, MN.

Cook, D. J., Cusimano, M. D., Tator, C. H., & Chipman, M. L. (2003). Evaluation of the thinkfirst canada, smart hockey, brain and spinal cord injury prevention video. *Injury Prevention, 9*(4), 361-366.

Csikszentmihalyi, M. (1992). *Flow: The classic work on how to achieve happiness*. New York, NY: Harper Perennial.

Curnow, W. J. (2003). The efficacy of bicycle helmets against brain injury. *Accident analysis and prevention, 35*, 287-292.

de Castell, S., & Jenson, J. (2003). Op-ed serious play. *Journal of Curriculum Studies, 35*(6), 649-665.

Deci, E., & Ryan, R. (1985). *Intrinsic motivation and self-determination in human behavior*. New York: Plenum.

Din, F. S., & Caleo, J. (2000, February 16-19). *Playing computer games versus better learning*. Paper presented at the Annual Conference of the Eastern Educational Research Association, Clearwater, FL.

Dobson, M., Kyrylov, V., & Kyrylova, T. (2004, August 16-18). *Decision training using agent-based business strategy games*. Paper presented at the Seventh IASTED International Conference on Computers and Advanced Technology in Education, Kauai, HI.

Education Arcade. (2005). History and direction. Retrieved October 23, 2005, from http://www.educationarcade.org/modules.php?op=modload&name=Sections&file=index&req=viewarticle&artid=3&page=1

Egenfeldt-Nielsen, S. (2005). *Beyond edutainment: Exploring the educational potential of computer games*. Unpublished Dissertation, IT-University of Copenhagen, Copenhagen.

Engeström, Y. (2001). Expansive learning at work: Toward an activity theoretical reconceptualization. *Journal of Education and Work, 14*(1), 133-156.

ESA. (2005). Essential facts about the computer and video game industry. Washington, D.C.: Entertainment Software Association.

Fabricatore, C. (2000). Learning and videogames: An unexploited synergy, *AECT National Convention*. Long Beach, CA.

Faria, A. J., & Wellington, W. J. (2004). A survey of simulation game users, former-users, and never-users. *Simulation & Gaming, 35*, 178-207.

Feinstein, A. H., & Cannon, H. M. (2002). Constructs of simulation evaluation. *Simulation & Gaming, 33*(4), 425-440.

Foreman, J. (2003, July/August). Next-generation educational technology versus the lecture. *EDUCAUSE Review, 38*, 12-22.

Games-to-Teach. (2005). Games-to-teach homepage. Retrieved October 24, 2005, 2005, from http://www.educationarcade.org/gtt/

Garris, R., Ahlers, R., & Driskell, J. E. (2002). Games, motivation, and learning: A research and practice model. *Simulation & Gaming, 33*(4), 441-467.

Gee, J. P. (2003). *What video games have to teach us about learning and literacy*. New York: Palgrave Macmillan.

Haddon, L. (1988). Electronic and computer games: The history of an interactive medium. *Screen, 29*(2), 52-73.

Hogle, J. G. (1996). *Considering games as cognitive tools: In search of effective "edutainment"* (No. ED 425 737): University of Georgia.

Jacobs, J. W., & Dempsey, J. V. (1993). Simulation and gaming: Fidelity, feedback and motivation. In J. V. Dempsey & G. C. Sales (Eds.), *Interactive instruction and feedback*. Englewood Cliffs, NJ: Educational Technology.

Jansz, J., & Martens, L. (2005). Gaming at a lan event: The social context of playing video games. *New Media & Society, 7*(3), 333-355.

Jonassen, D. (1997). Instructional design model for well-structured and ill-structured problem-solving learning outcomes. *Educational Technology Research and Development, 45*(1), 65-94.

Jonassen, D. (1999). Designing constructivist learning environments. In C. M. Reigeluth (Ed.), *Instructional design theories and models: A new paradigm of instructional theory* (Vol. II, pp. 215-239). Mahwah, NJ: Lawrence Erlbaum Associates.

Kirriemuir, J. (2003). *The relevance of video games and gaming consoles to the higher and further education learning experience* (No. TSW 02-01): Joint Information Systems Committee.

Kirriemuir, J., & McFarlane, A. (2003). *Literature review in games and learning*: NESTA Futurelab.

Kyrylov, V., Bonanni, C., Kyrylova, T., & Love, E. (2004). Advanced simulation gaming technology and business education: Simon Fraser University.

Lancaster, H. (1998, October 6). Today's free agents work at twitch speed: A buzzword glossary. *Wall Street Journal*, p. B1.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, England: Cambridge University Press.

Leemkuil, H., de Jong, T., de Hoog, R., & Christoph, N. (2003). Km quest: A collaborative internet-based simulation game. *Simulation & Gaming, 34*(1), 89-111.

Leyland, B. (1996). *How can computer games offer deep learning and still be fun? A progress report on a game in development.* Paper presented at the Australian Society for Computers in Tertiary Education (ASCILITE), Adelaide, Australia.

Lieberman, D. A. (2001). Management of chronic pediatric diseases with interactive health games: Theory and research findings. *Journal of Ambulatory Care Management, 24*(1), 26-38.

Ludology. (2005). Research centers. Retrieved November 18, 2005, from http://ludology.org/links.php?category=Research+centers

Macedonia, M. (2002). Games soldiers play. *IEEE Spectrum, 39*(3), 32-37.

Malone, T. W. (1981). Toward a theory of intrinsically motivating instruction. *Cognitive science, 5*, 333-369.

Martinez-Gallagher, M. M., & Norton, D. (2004). Outbreak quest: A 90-day game development initiative. In A. A. D. L. Co-Lab (Ed.): University of Wisconsin.

Maushak, N. J., Chen, H.-H., & Lai, H.-S. (2001). *Utilizing edutainment to actively engage k-12 learners and promote students' learning: An emergent phenomenon.* Paper presented at the National Convention of the Association for Educational Communications and Technology, Atlanta, GA.

McCrory, P. (2001). Do mouthguards prevent concussion? *British Journal of Sports Medicine, 35*, 81-82.

McCrory, P., Collie, A., Anderson, V., & Davis, G. (2004). Can we manage sport related concussion in children the same as in adults? *British Journal of Sports Medicine, 38*, 516-519.

McIntosh, A. S., & McCrory, P. (2005). Preventing head and neck injury. *British Journal of Sports Medicine, 39*, 314-318.

McMullen, D. W. (1987). *Drills vs. Games--any differences? A pilot study* (No. ED335355).

Miller, C. (2000). Designing for kids:Infusions of life, kisses of death. Retrieved October 23, 2005, 2005, from http://www.gamasutra.com/features/20000112/kids_01.htm

Mitchell, A., & Savill-Smith, C. (2004). *The use of computer and video games for learning.* London: Learning and Skills Development Agency.

Moon, J. (1999). *Reflection in learning & professional development: Theory and practice.* London: Kogan Page.

Murray, J. (1997). *Hamlet on the holodeck: The future of narrative in cyberspace.* New York: The Free Press.

Muzzy Lane. (2005). Company. Retrieved November 21, 2005, from http://www.muzzylane.com/company/

Nielson, J. (1993). *Usability engineering.* Boston: Academic Press.

NSF. (2005). Games and education grants. Retrieved October 24, 2005, from http://www.nsf.gov/awardsearch/piSearch.do?SearchType=piSearch&page=1&QueryText=games+education&PIFirstName=&PILastName=&PIInstitution=&PIState=&PIZip=&PICountry=&Search=Search#results

Okan, Z. (2003). Edutainment: Is learning at risk? *British Journal of Educational Technology, 34*(3), 255-264.

Perkins, D. N. (1991). Technology meets constructivism: Do they make a marriage? *Educational Technology, 31*(5), 18-23.

Prensky, M. (2001). *Digital game-based learning*. New York: McGraw-Hill.

Reber, A. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology, 118*(3), 219-235.

Reigeluth, C. M., & Schwartz, E. (1989). An instructional theory for the design of computer-based simulations. *Journal of computer-based instruction, 16*(1), 1-10.

Richtel, M. (2005, August 22, 2005). Once a booming market, educational software for the pc takes a nose dive. *The New York Times*, p. C1.

Rieber, L. P. (1996). Seriously considering play: Designing interactive learning environments based on the blending of microworlds, simulations, and games. *Educational Technology Research and Development, 44*(2), 43-58.

Robinett, W., & Grimm, L. (2005). Rocky's boots. Retrieved October 23, 2005, from http://www.warrenrobinett.com/rockysboots/

Rose, M., Haider, H., & Buchel, C. (2005). Unconscious detection of implicit expectancies. *Journal of Cognitive Neuroscience, 17*(6), 918-927.

Saettler, L. P. (1968). *A history of instructional technology*. New York: McGraw-Hill.

Salen, K., & Zimmerman, E. (2004). *Rules of play: Game design fundamentals*. Cambridge, Massachusetts: The MIT Press.

Sauvé, L., Renaud, L., & Kaufman, D. (2005, June 16-20). *Games and simulations: Theoretical underpinnings*. Paper presented at the DiGRA 2005. Changing Views: Worlds in Play, Vancouver, B.C.

Serious Games. (2005). About. Retrieved October 23, 2005, from http://www.seriousgames.org/about.html

Shanks, D., Rowland, L., & Ranger, M. (2005). Attentional load and implicit sequence learning. *Psychological Research, 69*, 369-382.

Silverman, B. G., Holmes, J., Kimmel, S., & Branas, C. (2002). Computer games may be good for your health. *Journal of Healthcare Information Management, 16*(2), 80-85.

Skinner, B. F. (1948). "superstition" in the pigeon. *Journal of Experimental Psychology, 38*, 168-172.

Squire, K. (2002). Cultural framing of computer/video games. *Game Studies, 2*(1).

Squire, K. (2004). *Replaying history: Learning world history through playing civilization iii*. Unpublished Dissertation, University of Indiana, Bloomington, IN.

Squire, K., & Jenkins, H. (2004). Harnessing the power of games in education. *Insight, 3*(1), 5-33.

Squire, K., Jenkins, H., Holland, W., Miller, H., O'Driscoll, A., Tan, K. P., et al. (2003). Design principles of next-generation digital gaming for education. *Educational Technology, 43*(5), 17.

SSHRC. (2005). Studies to probe innovative uses of technology in health care. Retrieved October 24, 2005, from http://www.sshrc.ca/web/whatsnew/press_releases/2003/ine_cura_sfu_e.asp

Steinkuehler, C. (2004). *Learning in massively multi-player online games.* Paper presented at the Sixth International Conference on Learning Sciences, Los Angeles, CA.

Svane, T., Aderklou, C., Fritzdort, L., & Hamilton-Jones, J. (2001). *Knowledge by user demand and self-reflection: New models for teaching and assessment in edutainment software design.* Paper presented at the 31st ASEE/IEEE Frontiers in Education Conference, Reno, NV.

Thiagarajan, S. (1998). The myths and realities of simulations in performance technology. *Educational Technology, 38*(5), 35-41.

van der Mast, C. (1995). *Developing educational software: Integrating disciplines and media.* Unpublished Doctoral dissertation, Delft University of Technology, Delft.

van Merriënboer, J. J. G. (1997). *Training complex cognitive skills: A four-component instructional design model for technical training.* Englewood Cliffs, NJ: Educational Technology.

Williamson, I., & Goodman, D. (In press). Converging evidence for the under-reporting of concussion in youth ice hockey. *British Journal of Sports Medicine.*

Wolfe, J., & Crookall, D. (1998). Developing a scientific knowledge of simulation/gaming. *Simulation & Gaming, 29*, 7-19.

Zipf, G. K. (1949). *Human behavior and the principle of least effort.* New York: Addison-Wesley.