Bridging the Lab and the Field: Three Essays on Experiments in Virtual Worlds

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

Peter Twieg
Master of Science
George Mason University, 2009
Bachelor of Science
Reed College, 2006

Director: Dr. Kevin McCabe, Professor
Department of Department of Economics

Spring Semester 2017
George Mason University
Fairfax, VA

Acknowledgments

I would firstly and foremostly like to thank my advisor, Dr. Kevin McCabe, for his support and mentorship throughout the creation of this dissertation. Along this journey I have also had the honor of meeting and conducting research with scholars such as Dr. Daniel Houser, Dr. Omar al-Ubaydli, Dr. David Eil, and Dr. Frank Krueger.

I owe a broad and deep debt of gratitude to all the individuals I have worked with at the Center for the Study of Neuroeconomics over the years. In particular, I would like to thank Kathy McCabe and Heather Leahy for providing continuous operational and organizational advice and support for my work. I would also like to thank Jaap Weel for first introducing me to Dr. McCabe as well as helping me develop the software skills that would eventually become critical to my research and career.

Table of Contents

List of Tables

List of Figures

Abstract

BRIDGING THE LAB AND THE FIELD: THREE ESSAYS ON EXPERIMENTS IN VIRTUAL WORLDS

Peter Twieg, PhD

George Mason University, 2017

Dissertation Director: Dr. Kevin McCabe

One persistant concern without experimental economics is the problem of generalizability from behaviors observed within the experimental laboratory to various field contexts. Described as the problem of "parallelism" by Vernon Smith, these issues have inspired the development of a variety of methodological tools within experimental economics. This dissertation focuses on one such tool that has risen to prominence in recent years: The use of online virtual worlds as that manage to blend the different advantages of laboratory research and field experiments in order to enable the exploration of new economic questions using experimental methods.

This first chapter of this dissertation represents an introduction and outline of the three essays that comprise its primary research content. The second chapter provides a discussion of the history of experimentation in virtual worlds and outlines some advice and guidelines for the construction and execution of these experiments. In the next two chapters I describe the results of two experiments that use virtual world environments in order to explore different questions concerning the establishment of informal property rights and institutions within groups that face different social dilemmas.

In Chapter 3, I explore the determinants of effective commons management institutions

and demonstrate that providing groups with opportunities for costly specialization generate stronger territoriality and better overall resource management. In Chapter 4, I explore how networks of trust and trade emerge amongst groups under conditions where opportunities for production, specialization and exchange are present but access to productive opportunities have a varying degree of scarcity and are governed by a system of squatters' rights. I show that as this scarcity increases, the efficiency of trade within groups declines.

Chapter 1: Introduction

In many traditional laboratory experiments conducted within the purview of experimental economics, subjects are recruited into a laboratory and presented with decision-making tasks that represent an abstracted version of some microeconomic system of real-world interest. In Smith (1982), Nobel Laureate Vernon Smith outlines the challenges of maintaining parallelism between controlled experiments and these real-world systems and the necessity of taking care to ensure that any results gleaned from the former will provide genuine insights about the latter.

However, guaranteeing parallelism from laboratory results is not an easy task. Because of these difficulties, experimentalists have explored a wide variety of methodological tools that make different tradeoffs between internal validity and external validity in exploring various research questions. This dissertation focuses on one such tool: The use of online virtual worlds to generate contextually-rich and naturalistic environments within which new microeconomic systems can be implemented and investigated by experimenters. Although virtual worlds as a technology have existed for decades, it has only been within the past ten years that experimental economists have begun exploring the potential for these platforms to serve as a methodological bridge between the carefully-controlled, replicable experiments performed in the laboratory and the high intrinsic parallelism of natural field experiments (Harrison and List, 2004).

There are several major goals of the essays that comprise this dissertation. First, I wish to convey to the reader a basic familiarity with the sorts of questions which virtual worlds have been used to answer and an appreciation for why virtual worlds were considered as an optimal platform for studying these questions. Second, I want to provide would-be experimenters with some useful tools and advice for getting started with the seemingly-daunting task of implementing a virtual world experiment - I believe that the unfamiliarity with the actual software tools that virtual world experiments tend to rely on is a major obstacle to their widespread adoption. Third, I aim to provide concrete examples of some of the virtual world experiments that I have performed and how they were used to contribute to our understanding of how humans establish informal institutions to solve social dilemmas.

The second chapter of this dissertation provides an introduction to virtual world research, outlining the advantages that virtual world designs can provide in terms of allowing experimenters to implement microeconomic systems with a deeper degree of control and parallelism than what can be afforded by alternative methodologies. Several different types of virtual world experiments are outlined before the narrative focuses in on the type of experiment that I'm most familiar with: Laboratory experiments using the OpenSim platform. Much of this chapter is devoted to providing general advice concerning OpenSim experimentation as well as outlining a general framework for the implementation of an economic experiment within OpenSim. A crucial goal of this chapter is to help experimenters overcome the fixed costs involved with becoming proficient in this software and how to use it in order to conduct experiments.

The last two chapters outline two specific experiments performed within virtual worlds. Both of these experiments explore questions of how groups establish and maintain informal institutions, leveraging the built-in spatiotemporal features and functionalities of OpenSim to explore their respective research questions. In Chapter 3, I explore the determinants of effective commons-management institutions and demonstrate that providing groups with opportunities for costly specialization generates stronger territoriality in resource foraging. Within this environment, groups are able to build these institutions and maintain them by employing several of the institutional design principles outlined in Ostrom (1990), including natural language communication, mutual monitoring and costly punishment. The spatial nature of this environment is crucial to the definition of territoriality in this environment: By borrowing a territorial measure used in ecological studies, I assign to subjects a harvesting region and examine the extent to which these regions overlap in different treatments.

In Chapter 4, I explore how networks of trust and trade emerge among groups under conditions where opportunities for production, specialization and exchange are present but access to productive opportunities may be scarce. In this environment, subjects utilize "harvesters" to procure primary goods that can be used in exchange, but the property rights that individuals can exercise over these harvesters is limited — if a subject leaves a harvester to find a trading partner that harvester can be claimed by another subject. The treatment conditions examined in this experiment modulate the number of harvesters active in each group, and when the number of group members exceeds the number of harvesters an excess supply of labor is generated. My results show that the interaction between weak property rights and the presence of excess labor has deleterious consequences for the effectiveness of exchange, resulting in an elevated amount of deadweight loss.

Chapter 2: An Experimentalist's Guide to Economic Experiments in Virtual Worlds

2.1 Motivating Virtual Worlds Research

In his work outlining the concept of an experimental microeconomic system, Smith (1982) emphasizes the attention that experimental economists should give to ensuring that their designs exhibit what he calls "parallelism" with behavior of interest in the field:

> Once replicable results have been documented in laboratory experiments, one's scientific curiosity naturally asks if these results also apply to other environments, particularly those of the field. Since economic theory has been inspired by field environments, we would like to know, if we were lucky enough to have a theory fail to be falsified in the laboratory, whether our good luck will also extend to the field. Even if our theories have been falsified, or if we have no theory of certain well-documented behavioral results in the laboratory, we would like to know if such results are transferable to field environments. (pp. 936)

Since this time, many advances in understanding the conditions under which parallelism holds stem from an improved understanding of behavior through the lens of what Smith (2003) labels "ecological rationality":

> Ecological rationality uses reason rational reconstructionto examine the behavior of individuals based on their experience and folk knowledge, who are "naive" in their ability to apply constructivist tools to the decisions they make; to understand the emergent order in human cultures; to discover the possible intelligence embodied in the rules, norms, and institutions of our cultural and biological

> heritage that are created from human interactions but not by deliberate human design. (pp. 470)

An early and well-documented example of the importance of understanding ecological rationality concerns the Wason selection task first defined in Wason (1968). In Wason's original investigation, subjects overwhelmingly failed to solve the deductive reasoning problem posed to them in the task. However, subsequent studies showed that these failure rates can be ameliorated when the reasoning tasks are posed in contextualized form that are similar to the sorts of problems that subjects face in their everyday lives (Cosmides and Tooby, 1992). Over the years, economists have investigated how framings and contextual cues can impact behavior within a variety of contexts (Kühberger, 1998; Lévy-Garboua et al., 2012; Liberman, Samuels and Ross, 2004; Tversky and Kahneman, 1985).

Experimentalists have also explored and identified several ways in which laboratory procedures may undermine parallelism, such as the flat maximum problem (Harrison, 1989) or experimenter demand effects (Bardsley, 2008; Zizzo, 2010). Levitt and List (2007) unifies many of these concerns, identifying five major possible challenges to external validity in laboratory procedures that can undermine parallelism between simple laboratory procedures meant to measure social preferences and field behavior:

- The presence of moral and ethical considerations.
- The nature and extent of the monitoring of decisions by others.
- The presence or absence of contextual cues regarding the decisionmaking task.
- Selection effects into experimental subject pools.
- The flax maximum problem.

One major tool for preserving parallelism and addressing possible issues with the external validity of laboratory tasks has been to find ways of implementing controlled randomizations in the field. In Harrison and List (2004), four major types of experiments involving greater degrees of naturalism are identified. Conventional lab experiments use standard subject pools (usually students) with abstract framing and rules. Artefacual field experiments use nonstandard subject pools of interest but preserve the abstraction of the decision task. Framed field experiments build on artefacual field experiments by introducing contextual cues and naturalistic framings. Lastly, natural field experiments take place in a framed environment where subjects would naturally engage in the experimental task and do not even know that they are participating in an experiment.

Despite the contributions of the growing field experiment literature to our understandings of important phenomena such as how poverty effects decisionmaking (Haushofer and Fehr, 2014; Mani et al., 2013) and how real-world groups utilize pre-existing institutions and social capital to solve various social dilemmas[1], more-natural field experiments have still received their own sets of methodological criticisms. They are often costly to perform and difficult to replicate. Furthermore, it is often difficult for experimenters to maintain a controlled environment and subject contamination is a constant concern (Banerjee and Duflo, 2009; Camerer, 2011). Relatedly, because experimenters often cannot directly observe, record, or manipulate all facets of subject decisionmaking and communication in natural field experiments, some variables of interest may either be unmeasurable or subject to bias in their measurement (Blattman et al., 2016).

Because of these limitations, experimentalists have explored a variety of tools that can be used to try to bridge the laboratory and the field, melding the advantages of field experiments that use naturalistic contexts and nonstandard subject pools while also retaining the high degree of experimenter control and reproducibility associated with traditional laboratory experiments. One important technology that is being adopted gradually by experimenters is the use of context-rich virtual environments that aim to replicate many of the salient features of different field environments. These virtual environments may come in many forms, but for general purposes a virtual environment will be defined in this paper

[1]For example, Oosterbeek, Sloof and Van De Kuilen (2004) provides a summary of the substantial cross-cultural variation in ultimatum game behavior. Additionally, Cárdenas and Ostrom (2004) and Prediger, Vollan and Frölich (2011) provide examples of experiments showing group composition and cultural influences in commons management tasks.

as a subject interface that possesses the following three features:[2]

- *Graphical Interfaces*: Features of the relevant economic environment are conveyed to the subject largely through graphical representations. Furthermore, the subject should have some control over this representation in order to induce social presence in the virtual environment.

- *Naturalistic Spatiotemporal Dynamics*: The environment is represented as being both spatial and temporal in nature. Subjects will interact with objects (and possibly one another) within this space over time. These dynamics do not have to wholly based on real-world settings, but more realistic dynamics are likely to exhibit greater parallelism.

- *Contextualization*: The environment utilizes some sort of naturalistic framing to contextualize its underlying structure. This framing will often relate to the real-world problems that motivate the underlying research questions.

A couple examples of virtual environments used in laboratory experiments are shown in Figures 2.1 and 2.2. Figure 2.1 depicts the environment used in Fiore et al. (2009), where subjects engaged in risky decisionmaking in a task concerning protecting one's property from brush fires. The authors demonstrated that using this naturalistic virtual environment to frame these choices generated more-accurate assessments of risk than a simple text-based interface. Figure 2.2 depicts the environment used in DeScioli and Wilson (2011), wherein subjects controlled avatars and used the graphical interface to engage in various forms of social interaction.

These studies have utilized virtual environments largely because they allow for features to be implemented that may enhance parallelism. One important advantage of virtual environments is that they induce a sense of social presence that may be absent in conventional laboratory settings (Fiore et al., 2009; Slater, Usoh and Steed, 1994; Spann et al., n.d.).

[2] A separate typology of virtual environments can be found in Innocenti (2015).

Figure 2.1: Fiore et al. (2009)

Additionally, avatar-to-avatar interactions within the virtual world can have beneficial effects on external validity by allowing experimenters to naturalistically reduce or modulate the social distance between subjects (Fiedler, Haruvy and Li, 2011; Greiner, Caravella and Roth, 2014) In their study on risky decisionmaking, Fiore et al. (2009) offer the following explanation of why they opted to use a “VX” (“virtual experiment”) environment:

> The VX environment will generate internal validity since it is able to closely mimic explicit and implicit assumptions of theoretical models, and thus provide tight tests of theory; it is also able to replicate conditions in past experiments for robustness tests of auxiliary assumptions or empirically generated hypotheses. The VX environment will generate external validity because observations will be made in an environment with cues mimicking those occurring in the field. In addition, any dynamic scenarios can be presented in a realistic and physically consistent manner, making the interaction seem natural for the participant. Thus the VX builds a bridge between the lab and the field, allowing

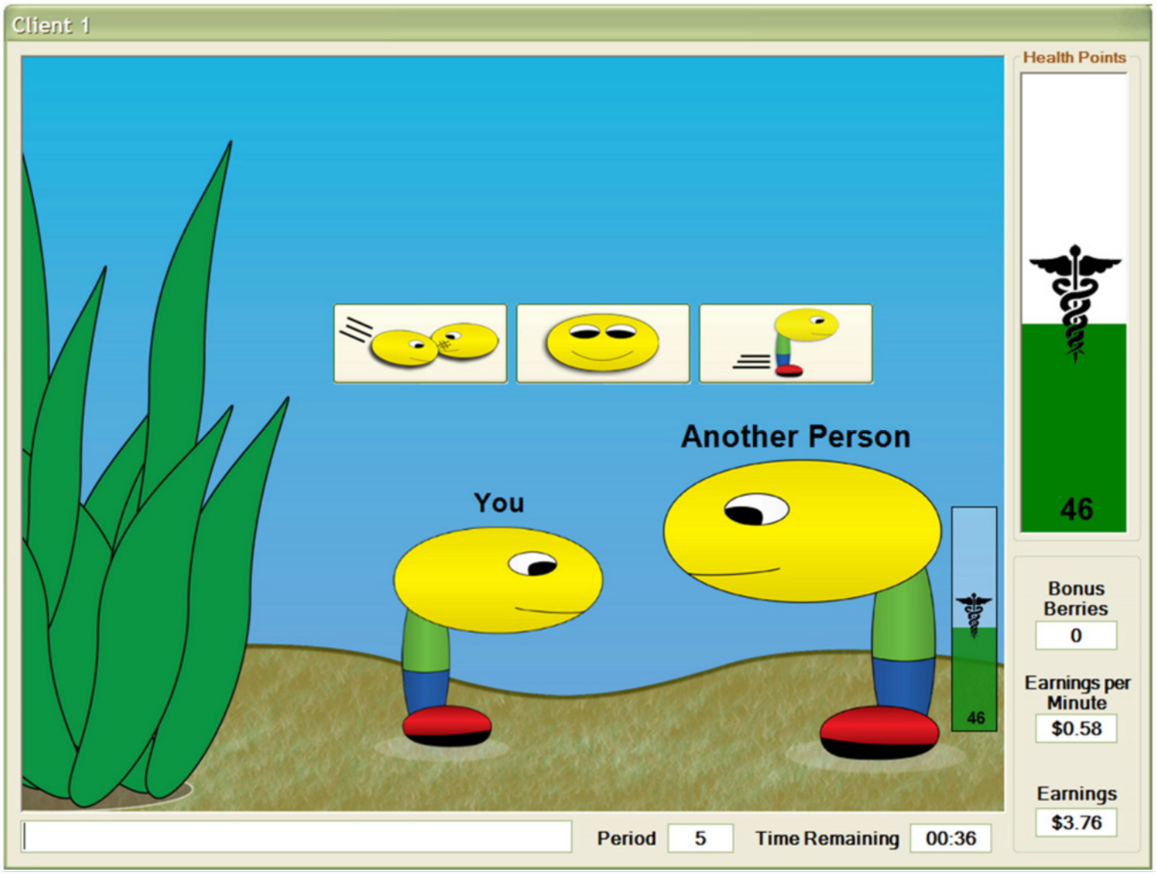

Figure 2.2: DeScioli and Wilson (2011)

> the researcher to smoothly go from one to the other and see what features of each change behavior. (pp. 66)

Elinor Ostrom has taken this line of argument further in her research, arguing that simulating the spatial and temporal dynamics that are involved in collective action problems such as commons management is crucial to understanding how groups in the field organize to address these dilemmas. In Janssen and Ostrom (2008), the authors observe the key roles that these dynamics play in the development of informal institutions:

> The participants in these experiments share a common renewable resource that is spatially explicit. We allow the participants to communicate about how they might improve their performance in harvesting the resource (in other words, how

> they can earn more money in the experiment). We will see that face-to-face communication enables them to discuss the common problem they are facing and to explore the feasibility of diverse methods for solving collective-action problems. Given the structure of the experiment, we observe that the types of informal rules participants propose are largely based on the spatial allocation of territories for individual use of the resource. (pp. 374)

This result would have been unattainable without providing subjects with a microeconomic environment that allowed for the observed institutions to be implemented by groups. The virtual setting of this experiment played a key role in providing subjects with the spatiotemporal dynamics that were a prerequisite for these institutions to be exist.

2.1.1 Virtual Worlds

The examples of virtual environments mentioned thus far have all been implemented via software that has been developed and customized for the research question at hand. However, there is another important class of virtual environments that have captured the interest of researchers: Virtual worlds. In general, these virtual worlds are persistent, avatar-based online platforms that contain a pre-existing pool of users. These virtual worlds often come with pre-defined contextualizations that experimenters can leverage to ask research questions. For example, Nicklisch and Salz (2008) (henceforth NS) use the virtual world of World of Warcraft in order to investigate a gift exchange model of labor supply. In this virtual world, fishing is a built-in economic activity that users can engage in, so the experimenters implemented a design that modulated how the conditions under which wages were paid to users (paid in the world's virtual currency) affected their productivity in this activity. In other virtual worlds, experimenters have significant content creation abilities that they can leverage in order to create their own microeconomic systems within the confines of the virtual world. Atlas and Putterman (2011) created a subject laboratory within the virtual world of Second Life in order to implement a trust game variant for the virtual world's users. These experiments often take advantage of a relatively cheap labor supply within

Figure 2.3: Avatar-based interface from Trust Networks

many virtual economies in order to cheaply generate a large number of observations.

My personal research has largely focused on developing experiments that take advantage of the built-in features of virtual worlds to answer questions concerning how groups establish institutions in order to overcome social dilemmas. For example, in the Trust Networks experiment outlined in Chapter 4 of this dissertation, subject groups are divided into different “villages” whose members can form intravillage and intervillage networks of trust and exchange. The spatial environment in this design serves to provide a naturalistic form of transaction cost to forming network links subjects will often have to choose between engaging in production at one location and finding trade partners at another. In the Berry Island experiment outlined in Chapter 3, subjects must navigate a spatial commons where resources can be freely foraged, but where overharvesting of resources can lead to a tragedy of the commons occurring. In this environment, the natural spatiality of the virtual setting allows for the investigation of how groups establish and maintain territorial property rights under different spatial distributions of resources.

The Berry Island design (which will be revisited throughout this paper as a motivating example) was inspired by the research of Marco Janssen and Elinor Ostrom (Janssen and Ostrom, 2008; Janssen et al., 2010) pertaining to the use of laboratory experiments to investigate commons management in spatial environments. This line of inquiry was inspired by the empirical finding that indigenous groups often defined informal usage rights along territorial lines (Acheson and Gardner, 2004; Brown and Pomeroy, 1999). In Janssen and Ostrom (2008), the authors assert that:

> A fourth key point that will be shown is that communication enables participants to find aspects of their own environment in the laboratory that are useful in making workable rules and monitoring each others conformance. Researchers have found that users of a common-pool resource tend to identify prominent aspects of the resource they are using that make it easier to monitor and regulate use patterns. Sometimes they create a spatial map using specific landmarks that are very obvious as ways of specifying territorial boundaries rather than some arbitrary, neat, rectangular array imposed by survey markers (Berkes 1986; Cordell 1984). Sometimes the natural layout of a territory makes entering a particular resource at a specific location more obvious and easier to monitor than others (Janssen and Ostrom 2006). (pp. 375)

Considerations like this played a large part in motivating me to extend this research into contextualized virtual worlds with field-like spatiotemporal dynamics. The particular virtual world that the Berry Island experiment was built in not only was well-suited for the construction of such an environment, but also made it easy provide subjects with basic tools such as mutual monitoring and natural language communication in order for subjects

Table 2.1: Field Context Factors of Two Virtual World Experiments

Field Context Factor	**NS (2008) Implementation**	**Berry Island Implementation**
Nature of subject pool	Subjects recruited by experimenter in virtual world, told they would be paid for performing a task.	Subjects recruited to a standard experimental lab and trained to use a virtual world.
Nature of subject information	Subjects already had a baseline knowledge of how to engage in the fishing task based on their prior virtual world experience.	Subjects had little or no experience with the environment or institutions used in this task.
Nature of the commodity	Subjects exchanged contextualized virtual resources with some intrinsic inworld value for currency.	Subjects collected a contextualized resource that they knew was depletable - "berries".
Nature of the task / trading rules	Defined through pre-existing rules governing fishing and exchange in World of Warcraft.	Given through the microeconomic system crafted by the experimenter.
Nature of the stakes	Subjects received inworld currency in exchange for fishing product.	Subjects received USD based on experimental dollars earned in the resource collection task..
Subject Environment	Used naturalistic fishing setting that already existed inworld.	Experimenter-constructed spatiotemporal environment where subjects explored a naturalistic setting and collected berries.

to establish a wide variety of rules and institutions for informal commons management in this environment. This environment not only enabled me as a researcher to develop a microeconomic system wherein territoriality could be instantiated, but to do so in a contextualized manner that should enhance the external validity of the findings of this research.

Table 2.1 shows how both the NS and Berry Island experiments line up as field experiments in the typology of Harrison and List (2004), who specify six major factors that define the field context of an experiment. Under this typology, experiments like NS that use a pre-existing virtual world's institutions and population would constitute a framed field experiment. Since Berry Island uses a traditional laboratory subject pool, it would not count as a field experiment using this particular hierarchy; however, it possesses field-like qualities on many of these factors due to its naturalistic contextualization.

At this point, one might be inclined to wonder why virtual environments aren't more commonly used in experimental research if they have a high level of potential.[3] The remainder of this paper will be dedicated to trying to address one of the major obstacles of conducting research using virtual environments: The high learning curve involved with mastering the software tools necessary to conduct this research. There is no equivalent to zTree (Fischbacher, 2007) for researchers interested in virtual environments, and most studies that uses them require an extensive amount of custom coding. However, once the fixed costs of mastering a tool have been shouldered, the marginal costs of creating rich and novel virtual environments to investigate new research questions can become mercifully low. A major aim of this paper is to convey information and insights that will lower these fixed costs for others who wish to conduct experimental research using virtual worlds.

2.2 Designing Virtual World Experiments

When it comes to implementing an experiment using a virtual world, it is crucial to understand how the given world's innate structure affects the design space available to experimenters. The presence of these structures will make implementing many environments and institutions much easier, but at the same time they may also serve to constrain the scope of microeconomic systems that may be investigated. In this section I will transition from discussing the broad conditions under which virtual world experiments may be desirable and instead focus on what the implications of these structures and constraints are for virtual world experimentation.

As an illustration, one major type of virtual world which exists in many forms is the Massive Multiplayer Online Role-Playing Game (MMORPG). MMORPGs generally have highly-contextualized worlds inspired particular intellectual properties, such as *Star Wars*,

[3]For recent reviews of published virtual world experiments, see Innocenti (2015), Gürerk et al. (2014), or Harrison, Haruvy and Rutström (2011).

Figure 2.4: The virtual world used in Nicklisch and Salz (2008), *World of Warcraft*.

Final Fantasy, or *Warcraft*. These virtual worlds tend to be highly structured by their designers and offer little leeway for experimentalists to construct their own environments. In particular, the NS experiment described earlier and depicted in Figure 2.4 was implemented in World of Warcraft. In this design, subjects were paid different fixed wages by experimenters for engaging in an inworld fishing activity. The reason why this fishing task was used was because this is one of the easiest labor tasks to have a subject perform within the given virtual world. Still, this task had some unavoidable features that could be considered undesirable from the perspective of an experimenter.

For example, in this experiment the labor efforts of subjects could not be directly observed, but instead used their observable fishing yield as a proxy for labor effort. However, it is a basic feature of *World of Warcraft* that fishing yield is actually a stochastic function of labor effort, so the number of fish obtained by subjects actually provided a noisy signal for labor effort. While it is unlikely that this effect compromises the authors' results, it illustrates how the underlying microeconomic experiment that was implemented was constrained by the built-in rules of the virtual world. There was no simple way for the experimenters to measure labor effort directly within *World of Warcraft*, and there was no way to make the proxy they relied on less-noisy because they did not have the ability to change the rules of the virtual world for their purposes.

One of the first choices that an experimentalist will be faced with when it comes to designing a virtual world experiment is choosing a platform within which the experiment will be implemented. The previous example is meant to highlight how the authors in NS were constrained in some regards with their experiment's implementation because of their decision to use *World of Warcraft*. While *World of Warcraft* could still provide experimenters with a platform within which a gift exchange model could be tested, there are many other microeconomic systems that would be infeasible to implement as controlled experiments within this virtual world. However, pre-existing virtual environments like *World of Warcraft* also provide opportunities for natural field experiments to occur using native subject pools. These subjects will often have a pre-existing understanding of the contextualized tasks, rules, and commodities that pertain to the experiment.

In contrast to using an MMORPG, the Berry Island and Trust Networks experiments described in the previous session were implemented using virtual worlds based upon a software platform called OpenSim. What is noteworthy about OpenSim-based virtual worlds, and what lends it much of its potential as an experimental platform, is the extent to which its users can generate and control their own environments and institutions through built-in content creation tools. OpenSim's virtual space is divided up into what are called "regions", where the vast majority of regions are owned by individual users (such as would-be experimenters) who have either rented the necessary server space and bandwidth to host these regions on existing virtual worlds or who have set up their own virtual worlds. The environment of each region, then, is defined by its owner it could be used for commercial purposes, or as a social destination, or it could be used to provide entertainment, or to fulfill any number of the owner's objectives. Each region can also be given its own individual theme and aesthetics one region might host a futuristic cityscape, while another might represent a peaceful forest, or a relaxing beach. Additionally, region owners are able to define who is able to access their environment they may choose to limit access to only certain individuals, or may grant free access to the public.

In contrast to using an MMORPG, the Berry Island and Trust Networks experiments described in the previous section were implemented within virtual worlds based upon a software platform called OpenSim. What is noteworthy about OpenSim-based virtual worlds, and what lends it much of its potential as an experimental platform, is the extent to which its users can generate and control their own environments and institutions through built-in content creation tools. OpenSim's virtual space is divided up into what are called regions, wherein the vast majority of regions are owned by individual users (such as would-be experimenters) who have either rented the necessary server space and bandwidth to host these regions on existing virtual worlds or who have set up their own virtual worlds. The environment of each region, then, is defined by its owner it could be used for commercial purposes, or as a social destination, or for the provision of entertainment, or any number of other objectives. Each region can also be given its own individual theme and aesthetics one region might host a futuristic cityscape, while another might represent a peaceful forest, or a relaxing beach. Additionally, region owners are able to define who is able to access their environment they may choose to limit access to only certain individuals, or may grant free access to the public.

This flexibility in content creation is the key reason why I have opted to use OpenSim as a platform for my research. The microeconomic environments underlying designs like Berry Island could not have been implemented (or roughly implemented) in a traditional MMORPG. To understand why this is, it is useful to discuss the different elements of a microeconomic system as defined in Smith (1982) and the considerations that an experimenter should keep in mind regarding them when defining a virtual world experiment that respects the advantages and disadvantages that different virtual worlds may provide.

2.2.1 Microeconomic Systems in Virtual Worlds

In Smith (1982), a microeconomic system is defined by two key elements: an environment and an institution. Each of the N agents in the environment, in turn, is defined via a utility function, a commodity vector, and a technological endowment. In virtual world experiments, experimenters will still have the freedom to shape utility functions by inducing value over subject actions and outcomes; however, they may face limitations in defining commodity spaces and technological endowments. To return to NS's use of *World of Warcraft*, the fishing design that was implemented utilized a set of commodities (fish) where there were already several environmental rules governing their acquisition and use. This saved the experimenters the trouble of having to define new commodities (which cannot be done in *World of Warcraft*), but also constrained them to have to use these environmental rules even when they may not be ideal as seen in the case of the stochastic production of fish.

Regarding institutions within a microeconomic system, Smith defines four critical components: A language consisting of a set of messages sent by agents, a set of allocation rules, a set of cost imputation rules, and a set of adjustment process rules. Attentiveness to the message space in particular is of critical importance to designing virtual world experiments. Most virtual worlds are avatar-based and provide some built-in functionality to allow users to communicate with one another, often using natural language. It is likely that experimenters will want to restrict the usage of these features in some ways. For example, in OpenSim, most software clients allow subject avatars to send private messages to one another. For the Berry Island and Trust Network experiments, however, this was considered undesirable for two reasons: Firstly, I wanted to be able to monitor all the messages sent and received by subjects during the experimental session, and these messages could not be

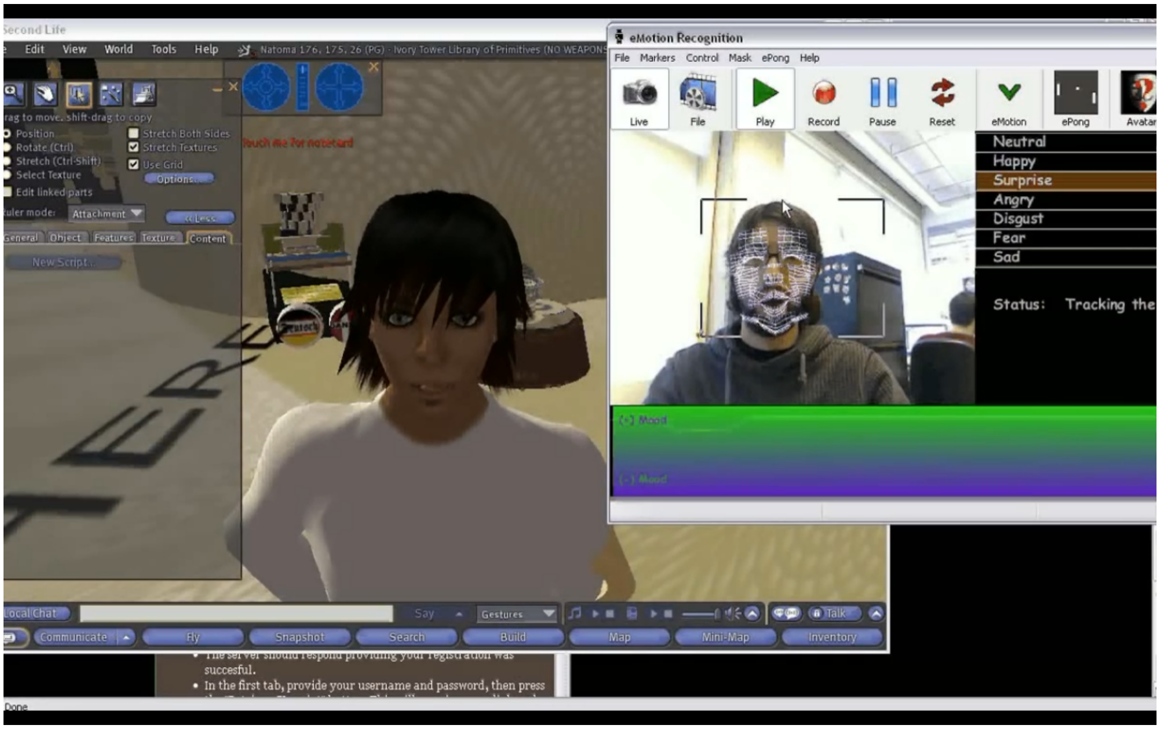

Figure 2.5: Facial tracking technology in Second Life.

easily observed and recorded by experimenters. Secondly, however, in many designs I simply didn't want subjects to be able to freely communicate. I wanted the spatiality inherent in these designs to provide naturalistic restrictions on how the avatars controlled by subjects could communicate two subjects should only be able to chat if they're in spatial proximity in the virtual world. These challenges were addressed through the use of the scripting tools available in OpenSim.

The message space of a virtual world will also contain several non-linguistic dimensions as well. For example, many virtual worlds allow for avatars to perform common gestures such as hand-waving that can be seen by other avatars. In my early virtual world experiments, I faced problems with avatars using tools to alter their own appearance during experimental sessions and had to specifically disallow this in the instructions before eventually figuring out how to disable this through scripting. The bottom line once again is that experimenters must be careful that these tools cannot be used by subjects in unforeseen ways that subvert the integrity of the underlying microeconomic system that the experimenter is seeking to implement. But as long as this is kept in mind, most virtual worlds will provide powerful software tools that will save experimenters the trouble of having to develop environments and institutions wholly from scratch. Many of these tools are on the cutting edge of current software technology. For example, researchers have developed software that allows for avatars in OpenSim and Second Life (a closely-related software platform which inspired the development of OpenSim) to reflect the facial expressions of users by using webcams, as depicted in 2.5. Additionally, several virtual worlds are committed to supporting immersion-enhancing technologies such as virtual reality headsets.

2.2.2 Control and Measurement in Virtual Worlds

By way of quoting Wilde (1981), Smith argues that:

> The fundamental objective behind a laboratory experiment in economics is to create a manageable "microeconomic environment in the laboratory where adequate control can be maintained and accurate measurement of relevant variables guaranteed" (Wilde, p. 138). "Control" and "measurement" are always matters of degree, but there can be no doubt that control and measurement can be and are much more precise in the laboratory experiment than in the field experiment or in a body of Department of Commerce data. (pp. 930)

Recall that a major goal of using virtual world experiments is to bridge the lab and the field by implementing experiments that have high parallelism while also preserving the critical advantages of the lab in terms of control and measurement. However, merely using a virtual world setting does not guarantee that control or the precise measurement of behavior can be taken for granted. A drawback of NS's design was that, similar to many field experiments, subject behavior was not directly observed. In *World of Warcraft*, monitoring subjects during the experiment would have required actually following them around with an experimenter avatar and making notes on their behavior. Not only would this be costly for experimenters to do, but this would be subject to human error and could potentially influencing subject decisionmaking by heightening the salience of the fact that their choices are being observed.

In general, experimenters aim to record as much information about subject decisionmaking as possible. In highly-structured virtual worlds like MMORPGs, this may be difficult. However, another reason I chose to use OpenSim was because as part of scripting microeconomic systems I could also ensure that the software recorded various decisions made by subjects. In the Berry Island experiment, I collected data such as the location of subjects in the spatial environment at periodic intervals, the chat messages that were sent to one another, and their harvesting decisions. Through the use of these scripts, I was able to record data at a level that would be very difficult to capture in the field.

As emphasized previously, it is critical to understand the virtual world in order to maintain control of an experiment. Each virtual world possesses a built-in action space that avatars will have access to, and it is the responsibility of the experimenter to make sure that either their microeconomic model accounts for these actions or that their use is restricted somehow. In OpenSim, for example, avatars are not only able to navigate spatial environments by walking, but are also allowed to fly and teleport from one location to another. For the sake of parallelism, these capabilities were disabled during my experiments. Furthermore, much like field experiments, virtual world experiments can be subject to various of contamination from outside factors. For example, if avatars from outside of the experiment are able to interact with subjects within the experimental setting, this could lead to various sorts of problems. Ultimately, what's important to take away from this discussion is that in many cases there is a tradeoff between taking advantage of the built-in features of a given virtual world and being constrained by how those features impact the sorts of microeconomic systems that can be implemented in the virtual world. It is crucial to understand and appreciate these tradeoffs when it comes to developing an experiment within a virtual world, and the microeconomic systems approach outlined in Smith (1982) provides a useful framework for doing so.

2.2.3 Microeconomic Systems in OpenSim

The software platform that I have used for most of my virtual world experiments, including Trust Networks and Berry Island, is OpenSim. At the moment there are several publicly-available virtual worlds based upon the OpenSim software, such as *Kitely* and *InWorldz*. Beyond these options, experimenters can set up and host their own virtual worlds in order to exercise greater control over the regions where their experiments take place. More detail concerning the considerations that go into setting up an OpenSim grid will be provided in Section IV.

The reason why OpenSim was chosen for my research was that it provided an optimal balance of built-in structure and flexibility in defining microeconomic systems of interest. OpenSim worlds are inherently spatial and avatar-based, and provide built-in support for various modes of interacting with both the environment (such as moving around, flying, and clicking on objects to interact with them in different ways) and with other avatars (such as communicating using typed messages and avatar gestures.) Built-in features that are often undesirable for experiments, such as the ability of avatars to freely teleport or fly or alter their appearance, can be easily disabled. Taken as a whole, these features provide a default environmental structure and message space to experimenters that can be altered to create a desired microeconomic system in a less-costly manner than having to create new software from scratch that implements these components. These environments are constructed through the built-in tools that are provided to alter OpenSim regions in order to shape interactive environments and avatar interfaces, as well as implement mechanisms for experimenters to monitor and record data from experimental sessions.

OpenSim, however, may not suit the needs of every experimenter interested in virtual worlds research. For example, NS managed to economize greatly on implementation costs by using World of Warcraft's built-in structures for their experiment. Furthermore, by using the native user popular of that virtual world, they were able to access a non-standard subject pool for their framed field experiment. However, World of Warcraft allows for very little flexibility for experimenters to do more than tweak an existing set of pre-defined microeconomic institutions. Because of this flexibility, OpenSim and the closely-related Second Life platform have become a common platform used in many virtual world experiments.[4]

In order to aid experimenters with implementing their own microeconomic experiments within virtual worlds, the remainder of this paper will focus on using the OpenSim platform to run experiments. This narrowing of scope is necessary going forward since describing the idiosyncratic tools and considerations necessary for running experiments in many different virtual worlds would require much more detail than what can be contained in a single paper, and in many cases these discussions would exceed the boundaries of my expertise. Nevertheless, many of the topics that will be covered as part of detailing the implementation of an OpenSim experiment will be applicable to the development and administration of experiments that use other virtual worlds.

2.3 Implementing an OpenSim Experiment

2.3.1 Motivating Example: The Tragedy of the Commons

The purpose of this section is to provide a first-pass understanding of how OpenSim environments can be constructed for use by experimenters. This will be done by way of example, where I will walk through the deployment of a simplified version of the Berry Island experiment referenced in the previous section. Along the way, the various considerations that go into each design choice will be explored. Going forward, this experiment will be referred to as the "Example Commons."

In the Example Commons, the environment consists of a set of "berry bushes" that are spatially distributed throughout the virtual world. This environment is susceptible to the tragedy of the commons due to the population dynamics of the berry bushes. Each berry bush can contain up to nine berries, and its population is governed by the logistic population growth model employed in Janssen et al. (2010). An exact specification of this

[4] For a non-exhaustive list of experiments that have employed OpenSim and Second Life, see Innocenti (2015).

model is formally defined in Chapter 3 when Berry Island is discussed in greater depth.The key population dynamics here pertain to the fact that when the population of a berry bush is below its carrying capacity of nine, then new berries will grow on that bush over time. Furthermore, the rate of berry harvesting is optimized when the berry bushes are kept at an intermediate population level that maximizes their overall rate of growth. If a bush is ever fully depleted, then its growth rate is 0 for the remainder of the experiment.

Subjects in this environment have access to a set of tools that can be used both to interact with the berry bushes as well as one another. Subjects who approach berry bushes can observe how many berries are present and harvest them in order to reduce this population. The tragedy of the commons arises in this experiment from the fact that the individual return to allowing the berry bushes to grow at an optimal rate rather than harvesting them immediately is less than the social return: In the absence of secure property rights any subject may come and harvest the berry bushes that others have been trying to cultivate. Ultimately, the central dilemma faced by subject groups is how to establish informal institutions that prevent the depletion of this common resource.

In order to overcome this dilemma, the Example Commons affords subjects several mechanisms that can be used for social interaction. Firstly, avatars can communicate with other nearby avatars using natural language. Secondly, since they share a spatial environment subjects have an ability to monitor nearby avatars and observe and react to the behavior of these avatars within the environment. Lastly, since costly punishment has been shown to be an effective tool for overcoming social dilemmas such as the tragedy of the commons (Carpenter, 2007; Denant-Boemont, Masclet and Noussair, 2007; Nikiforakis and Normann, 2008), subjects have the ability to punish in order to incentivize one another to not engage in undesired behavior. These abilities should jointly enable subjects to informally implement the sorts of institutional rules that predict successful commons management in field settings. Ultimately, each subject's payoff will be defined as follows:

$$U_i = B_i - \sum_{i \neq j} (p_{out} * P_{ij} + p_{in} * P_{ji}) \tag{2.1}$$

In equation 2.1, B_i represents the number of berries collected by subject i and P_{ij} represents the amount of punishment given from i to j, with the parameters p_{out} and p_{in} representing the costliness of both punishing and being punished.

In order to proceed in implementing this microeconomic system as an experiment, several design questions must be answered, including:

- How does one construct an interactive environment featuring berry bush objects and other contextual cues (geographical features, trees, etc.) that will be navigated by subject avatars?
- How does one build scripts so that berry bush populations reflect the dynamics of the underlying logistic growth function and update based on being harvested by avatars?
- How does one control what sorts of actions avatars can perform? As a baseline, subjects will have to be able to move around the environment and harvest berry bushes but how does one implement institutional tools such as natural language communication and costly punishment?
- How does one convey data about the state of a session to subjects? For example, if we want subjects to know how many berries they have harvested, how should this be done?
- How does the experimenter control the flow of the experiment? How is an experimental session initiated or halted? How can experimental parameters be changed in a flexible manner to allow for the smooth imposition of different treatments?
- How is data from the experiment recorded and organized for further analysis? What tools do experimenters have for observing this data in real time and noticing anomalies or bugs?

The remainder of this section will address these questions in turn. Additionally, a downloadable file containing a working version of the Example Commons along with the experiment's source code can be found online at `https://github.com/twiegp/example_commons`.

2.3.2 Building an Experimental Environment

The Example Commons experiment will be constructed and implemented within a single region of an OpenSim grid.[5] Furthermore, it will be assumed that the region will be owned by the experimenter. More details about how to acquire regions and the considerations that would favor using one type of grid over another will be given in Section IV, but properly weighing those considerations will require a familiarity with many of the concepts that will be discussed in this section. By default, each OpenSim region has a length and width of 256 meters, which takes an avatar roughly a minute to traverse at a normal running speed. While using larger environments may provide some upside in terms of naturalism (most real-world forests and fisheries, for example, encompass territories much larger than a sixteenth of a square kilometer), they also risk becoming unwieldly if subjects have to devote a large share of their time in an experiment simply to traversing the environment. As a consequence, larger-scale experiments that take place over longer time periods might benefit from using larger regions or multiple regions. Additionally, while it is possible to run experiments without owning a given region one can lease parcels of land from region-owners, as was done in Second Life by Atlas and Putterman (2011) this may come with limitations on one's ability to fully sculpt the environment as is necessary.

In OpenSim, most building and scripting is done while controlling an avatar inworld through the use of a "viewer", a client software package that communicates with the OpenSim server and renders the virtual world to users.[6] The fundamental building block for most

[5]The primary reason for this is that it is often easier to script objects that communicate data to one another when they are located on the same region.

[6]Further details about viewers and the different sorts of functionalies that they may have can be found in Section IV.

Figure 2.6: Building an inworld berry bush. The prim-editing interface is depicted on the right-hand side of the screen.

objects is called the prim (short for “primitive”)[7], and avatars can create prims within regions and manipulate them using tools that are built in to most viewers. Each prim has a basic shape, such as a sphere, cube, or cylinder, and can be modified to create detailed objects in the environment. Sets of prims can be linked together in order to create more complex objects. Figure 2.6 shows the rudimentary berry bush of the Example Commons experiment, which consists of two prims a rectangular green background prim and a red cylindrical “berry” prim with an attached texture meant to represent the number of berries available on that berry bush. The default forms of these two prims can be seen next to

[7]More recently, OpenSim has added support for mesh-based objects to be represented in world. This allows for the importation of mesh objects generated by software packages such as Blender to create highly-detailed inworld objects. For the purposes of subsequent discussion, mesh objects can be thought of as special single-prim objects.

the completed berry bush. Importantly, each prim in OpenSim has an inventory associated with it that allows for the prim to contain items such as scripts, textures, sound files, and even other objects comprised of prims. Managing this inventory is necessary in order to imbue these objects with functionality for example, if we want to make the numeric texture displayed on the berry part of the berry push to change with that bush's population, then the easiest way to achieve this would be to put both a set of texture objects in the berry's inventory as well as a script that will manage which texture is displayed.

The visual simplicity of the berry bush contrasts with the more detailed environmental objects that can also be seen in Figure 2.6, such as the gazebo seen behind it. This is a result of the fact that I did not personally design make these other objects, whose intricacy is beyond my ability to replicate. Instead, I acquired these assets as freebies made by other users and chose to use them to contextualize the environment of this experiment. This touches on one important feature of virtual worlds like Second Life and OpenSim each object, including our newly-created berry bush, has configurable permissions that allow for them to be transferred, copied, and bought or sold according to the wishes of the object's creator. As the creator of the berry bush object, I can freely make copies of it and decide whether subsequent owners can freely make copies as well. When it comes to building environments using multiple avatars, coordinating on object permissions is very important in order to allow other administrator avatars to manipulate the environment. But as shown in Figure 2.6, the fact that other content creators have made useful environmental objects that are freely copyable has saved me the trouble of having to make immersive environmental objects on my own.

These basic building features allow for the instantiation of the spatial environment that subjects will interact with during a session. In the Example Commons experiment, these objects consist of various contextual features (trees and grasses), berry bushes, and a set of walls around the border of the region to prevent subjects from leaving. Additionally, a campfire has been placed in the center of the region, which may provide a focal point for group coordination should subjects choose to meet around it.

```
// Bush configuration
integer INITIAL_BERRIES = 9;
float MAX_PICK_DISTANCE = 5.0;

// Bush state information
integer CURRENT_BERRIES;

default
{
    state_entry() {
        CURRENT_BERRIES = INITIAL_BERRIES;
        llSetTexture((string)CURRENT_BERRIES, ALL_SIDES);
    }

    // When touched by an avatar within range, if the bush has berries decrement the number shown on the texture.
    touch_start(integer num) {
        if (CURRENT_BERRIES > 0 && llVecDist(llDetectedPos(0), llGetPos())< MAX_PICK_DISTANCE) {
            CURRENT_BERRIES--;
            llSetTexture((string)CURRENT_BERRIES, ALL_SIDES);
        }
    }
}
```

Figure 2.7: A simple berry bush script written in OSSL.

2.3.3 Scripting an Experimental Environment

An Introduction to OSSL

The next step in building an OpenSim experiment is to imbue the objects in the environment with their desired functionalities. To accomplish this, OpenSim uses a C-like programming language called "OpenSim Scripting Language" (OSSL.) In order to be executed, each script must be placed within the inventory of a given prim. Figure 2.7 shows an example of a simple script that adds some basic functionality to our berry bushes when placed in the inventory of the berry prim upon initialization this prim will be set to display a texture indicating that the bush contains 9 berries, and each time it is clicked this texture display will decrement until there are no berries left on the bush.

This does not capture the full functionality that we wish to grant berry bushes in the Example Commons[8]; however, it serves as a starting point to introduce some important concepts regarding OSSL. OSSL is closely related to and derived from LSL (Linden Scripting Language), which is used to script objects in Second Life. A full description of how OSSL works and the various built-in functions that can be used in it is beyond the scope of

[8] Namely, the population growth dynamic does not exist in this script.

```
14     // When touched by an avatar within range, if the bush has berries decrement the number shown on the texture.
15     touch_start(integer num) {
16         if (CURRENT_BERRIES > 0 && llVecDist(llDetectedPos(0), llGetPos())< MAX_PICK_DISTANCE) {
17             CURRENT_BERRIES--;
18             llSetTexture((string)CURRENT_BERRIES, ALL_SIDES);
19             llRegionSay(-50,  llDetectedName(0) + ",HARVESTED");
20         }
21     }
```

Figure 2.8: A code excerpt from a berry script. Line 19 adds the ability for this object to communicate harvest events concerning the berry bush to other scripts.

this paper, although many useful resources for building objects using OSSL exist, such as Moore, Thome and Haigh (2009) and Heaton (2007). Linden Labs' LSL Portal[9] also provides constantly-updated information about various LSL functions and concepts which are largely applicable to OSSL as well.

Object Communication and HUDs

Let's say that we next want subjects to be able to track how many berries in total they've harvested during a session of the Example Commons experiment. Doing this involves solving two important problems: First, we must figure out a way to have the scripts that respond to harvesting events (in this case, the berry script) convey this information to users that is, when a user goes and clicks on a bush in the environment to collect a berry, how is information pertaining to this event conveyed to the user? Second, we must figure out a way to represent this information to a subject through their viewer. This section will address these two questions in turn.

One important feature of OSSL scripts is that they cannot directly access data from other scripts data can only be communicated between scripts in a limited number of ways. If our goal is to implement a script that tracks data on how many berries a user has collected across all berry bushes, then we will need to implement code that conveys information about harvesting events from the berry bushes being harvested to a script that's "listening" for these messages.

The bottom line is that almost all communication between prims in a Second Life

[9]http://wiki.secondlife.com/wiki/LSL_Portal

```
// Static global variables
integer HARVEST_VALUE = 2;

// Dynamic global variables
string USER;
integer HARVESTED;
integer PAYOFF;

// Function for updating the text display on the prim containing this script
updateText() {
    llSetText("
        Owned by " + USER + "
        Harvested Berries: " + (string)HARVESTED + "
        Current Payoff: " + (string)PAYOFF + "
    ", <1,1,1>, 1.0);
}

default
{
    state_entry() {
        // Initialize the script so it knows which user it is associated with.
        USER = llKey2Name(llGetOwner());

        // Define a listen event handler so it responds to messages.
        llListen(-50, "", NULL_KEY, "");

        // Initialize the text display on prim containing this script
        updateText();
    }

    listen(integer channel, string name, key id, string msg) {

        // First convert any string heard to a list, assuming that the first entry in the CSV-formatted list will be the intended recipient.
        list msgList = llCSV2List(msg);
        string recipient = llList2String(msgList,0);

        // If the recipient is the script's owner, then perform some action.
        if (recipient == USER) {
            string action = llList2String(msgList,1);

            if (action == "HARVESTED") {
                HARVESTED++;
                PAYOFF += HARVEST_VALUE;
                updateText();
            }
        }
    }
}
```

Figure 2.9: A user status script, version 1.0 This script listens for messages from the berry script in order to display information about a given subject's harvesting.

experiment will occur by sending text strings from one script to another in various ways. 2.8 excerpts a modified version of the berry script highlighted in Figure 2.7. When a user clicks on a berry bush containing this script, it will send a message out as a comma-separated string to any prim containing a script that is configured to listen to such a message. Figure 2.9 shows such an example of a message-receiving script. This script will listen for messages from other objects, and when a harvest message corresponding to a particular user (in this case, whoever owns the prim that the script is inside) is transmitted, the script will generate a text display above the prim displaying how many berries have been harvested by that user.

Figure 2.10: A close-up of a HUD running the user status script from Figure 2.9

How should the information in this script be conveyed to subjects? The best way to accomplish this is to design a Heads-Up Display (HUD) that contains the user status script. Like other objects, HUDs are comprised of prims which may have scripts inside, but what makes a HUD a HUD is that the object is attached to the user's field of vision so that the HUD stays in a fixed position on the screen as the subject's avatar navigates the environment.

Figure 2.10 shows a simple, 2-prim HUD where the user status script is placed in the bottom prim and text is displayed above it. Another major function of HUDs in experiments is to provide subjects with additional opportunities to interact both with their environment and with other avatars. For example, an interface for the punishment strategies that we wish to allow subjects to access can be implemented through the HUD. Let's say that we

Figure 2.11: The master object being used by an administrator to start a session.

want it to be the case that clicking on the gray prim on the bottom of the HUD executes a punishment event and deduct earnings from both the given subject and all other nearby avatars. This would be implemented by making it so that clicking on this prim would both update the score stored within the user's HUD script and also send a message out to other nearby HUDs that will be processed in order to deduct from those subjects' payoffs. While the code for this functionality is not shown here, it is implemented and documented in the full Example Commons experiment available at [url].

Controlling Experiment Flow

At this point several topics related to implementing features of the Example Commons environment have been covered, but no discussion has been made of how an administrator would actually start or stop or reset a given experimental session. To do this, we will need to create what will be called a Master object that contains scripts that allow the administrator[10] to perform these functions. Figure 2.11 shows an example Master object

[10] The administrator will normally be controlling an avatar on the same region as the experiment, but will be hidden from the view of the subjects one common way to achieve this is to simply place the Master

being used by an administrator's avatar. The HUDs and berry bushes will also be configured to receive and respond to messages from the Master object in order to initialize and halt in response to the appropriate messages.

One other important functionality of the Master object is that it provides a natural location for centralizing data about the current session. For example, let's say that in the Example Commons experiment we want to implement the following two treatment conditions:

- *Individual Growth (IG)*: The logistic growth of the number of berries on a given berry bush is a function of the population of berries only on that berry bush.
- *Group Growth (GG)*: The logistic growth of the number of berries on a given berry bush is a function of the population of berries across all berry bushes.

In the IG treatment, the experimenter could either make it so that each individual berry bushes performs periodic checks in order to determine whether its population should grow, or alternatively the Master object could perform an update for all the berry bushes at once and then communicate to them how many berries they should have. In the GG treatment, however, in order to perform an individual update a berry bush script needs to not only know how many berries it currently has, but it also needs to know how many berries exist across the entire experimental environment. While it could be the case that every berry bush communicates with every other berry bush in order to keep track of the state of the entire environment, there are efficiency gains to be had in terms of elegancy, processing time and memory by only having the Master perform these updates and then propogate out any population changes to the affected bushes.

With many of the relevant scripting concepts for OpenSim defined, we can now look at the communications blueprint of the entire Example Commons experiment shown in Figure 2.12. This blueprint has three primary components the environmental objects (berry bushes), the subject HUDs, and the Master object. When implementing a new

object very high in the sky where subject avatars will be unable to see it or reach it.

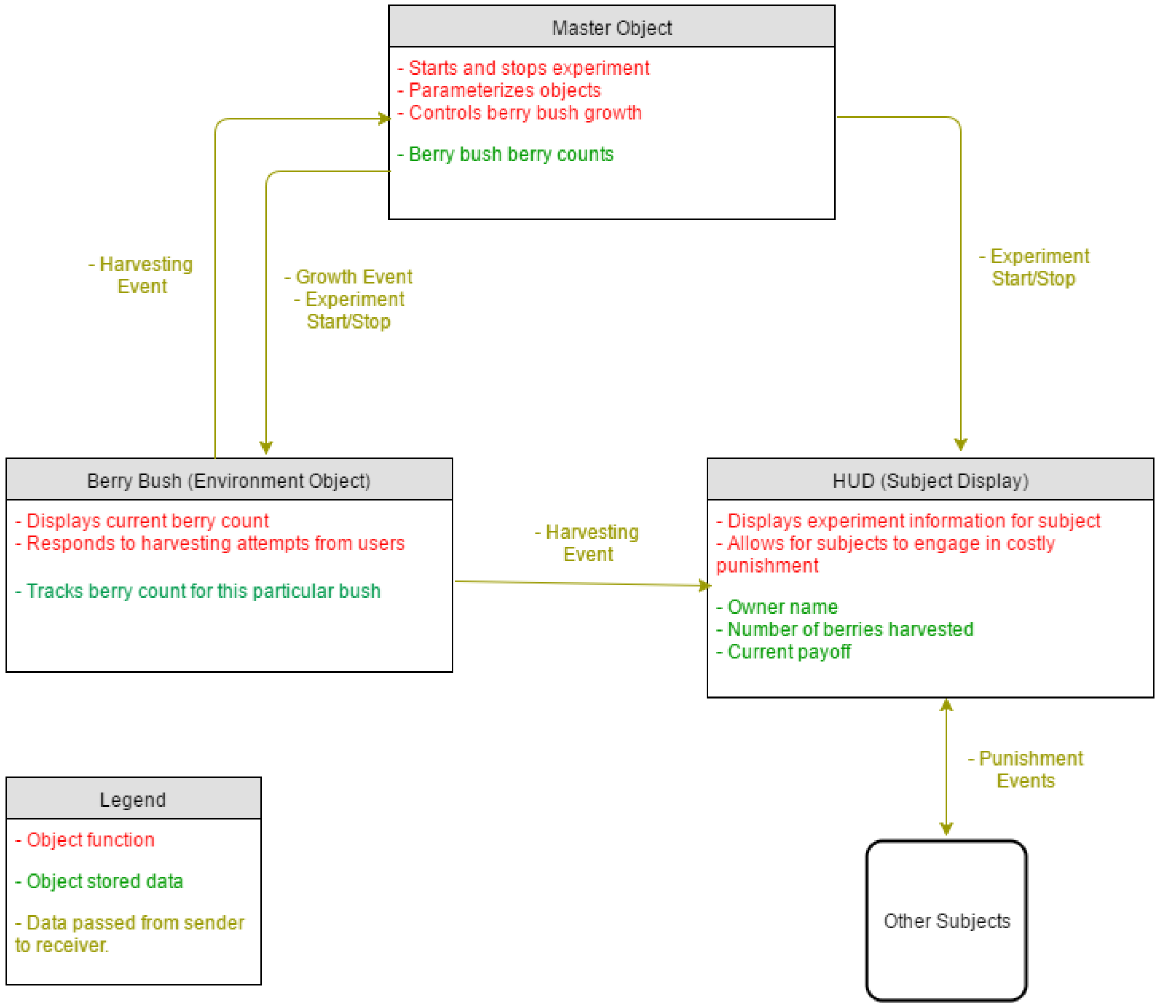

Figure 2.12: Depiction of full communication structure of the experiment, focusing on one berry bush and one subject.

OpenSim experiment, this is where I would recommend starting: Figure out what objects your experiment needs, and where various pieces of data will be stored and how how exactly each part will communicate with one another. With this expanded vocabulary, many of the design questions pertaining to the Example Commons outlined at the end of Section III.a can be made more particular:

- What objects need to exist in the environment? What data and functions will these objects need to store and execute locally?

- How should a HUD be designed to intuitively convey to subjects information that they should know during the session, and also to allow them to perform actions that aren't built-in to OpenSim, such as costly punishment?
- Beyond administering the experiment, what should the master object be doing? If any sort of information about the experiment needs to be centralized so that it's accessible by other objects, the Master object is the natural place to do this.

There are many subtle considerations that will affect how this blueprint will apply to a particular design. In general, it is wise to cut out redundant processing or storage of variables this is why having the Master instead of each berry bush keep track of the total number of berries in the environment was advocated in the previous section. If too much processing is going on during an experiment, it is possible that scripts might begin executing slowly, which could have a number of negative impacts on the experiment.

Data Collection

So far enough has been discussed to make the experiment fundamentally work from the perspective of subjects, but experimenters obviously need to extract data from subject behaviors to test their hypotheses. Unfortunately, in OpenSim one cannot simply do something like dump the state of a region periodically to a dataset. Since data is communicated between objects in OpenSim via string messages, data collection will likely consist of collecting the string messages generated by key events and then extracting them from OpenSim and transforming them into a flat dataset.

There are two general ways to record data generated in OpenSim. The easy way is to just have whatever messages you want to record relayed to the inworld admin avatar as chat messages[11], which the experimenter can then copy and paste to an external document or database at the conclusion of the session. The primary downside of this approach is that if the experimenter experiences any connectivity issues or simply forgets to record the

[11]This is most-readily done by having an object owned by the experimenter scripted to listen to the relevant messages and relay them to the experimenter via the llOwnerSay function.

messages this data might end up being lost in part or in whole.

In any case, what an experimenter will likely end up with is a chronologically-ordered set of messages. Some of these messages may correspond to particular events whose incidence was deemed worthy of record and some of them may correspond to periodic status updates that could be of use. Recording periodic updates could be useful for checking the internal consistency of the data, or for doing something like building a record of each avatar's location at periodic intervals over the session. It is likely that these event records will require significant post-processing to be transformed into an actual dataset, although there are various steps that can be taken to minimize the complexity of this task, such as standardizing the structure of different sorts of messages as much as possible.

Summary

At this point the questions laid out at the end of Section III.a have fully addressed. The Example Commons is comprised of three main objects a HUD for use by subjects, a set of berry bushes scattered around the environment, and a Master object used to start and stop sessions. Additionally, a number of free objects collected from other regions were used to contextualize the spatial environment.

Once the environment was made, OSSL scripts were defined to manage the functionality of each object and how they would need to communicate with one another to operate properly. By default in most OpenSim clients subjects are able to walk around and communicate with other nearby avatars using natural language, and this will be part of the Example Commons experiment. Additionally, we added the ability for users to punish all other members of their group by using the HUD. The HUD will also display for users information about their own personal harvesting and payoffs, as well as alerting them when they are punished. Furthermore, the experiment is parametrized through messages given from the master to the HUD and berry bushes at the beginning of the experiment. A proposed treatment variable Individual Growth vs. Group Growth can be toggled within the Master. An administrator can click on the Master object in order to use it to start and stop the experiment. During the session, data about subject decisions are relayed to the Master object and then to the administrator's avatar, which can then be collected for analysis.

The complete version of the Example Commons, available at [url], implements all of these features in the ways described. The various scripts used in the objects in the Example Commons are more-complicated than the ones used in the OSSL examples used earlier, but are well-documented and illustrate how all of these features are designed to work.

2.4 Useful Tools for OpenSim Experimenters

One of the goals of the previous section was to familiarize readers with many key principles of experiment design in OpenSim. With those concepts in mind, we can now discuss various tools that are available in the software ecosystem surrounding OpenSim and how they can be used in the broader process of experimental design and the effective implementation of microeconomic systems.

2.4.1 Choosing a Grid

One of the first choices that an experimenter must make is to decide on a grid where their experiment will be located. Recall that OpenSim is merely a software platform and that each individual virtual world that uses this platform comprises its own grid as such, all experiments in OpenSim are performed on grids. Different grids will vary on a number of dimensions, and the optimal bundle of features will depend on the needs of the experiment being implemented. For example, the primary Second Life grid that is owned and run by Linden Labs has the largest active userbase of any grid, and is relatively stable in terms of being bug-free and having a high uptime, but is also restrictive in terms of the degree to which experimenters can control regions. For example, one common difficulty that I encountered when running experiments on the Second Life grid was that there was a limit to the rate that individual scripts could send messages to other scripts, which necessitated some convoluted workarounds in order to ensure that experiments performed properly and that data was being correctly recorded. In order to help experimenters avoid these frustrations, I would advise that the following factors be considered when choosing a grid:

- **Native populations:** If one wishes to use the digital natives of a grid in experiments, then it is important to choose a relatively populous grid where recruitment costs can be minimized. The main Second Life grid is particularly attractive in this regard, but many alternative OpenSim grids also have sizable populations that can be reached through community hubs.

- **Administrative Control:** This issue can arise in many forms, but the general idea here is that some grids are more restrictive in how experimenters can manage regions than others. Oftentimes these restrictions exist in order to prevent malicious practices such as abusing users or utilizing excessive amounts of computational resources. The primary Second Life grid is especially restrictive in this regard, and many of these restrictions are poorly-documented, which can cause frustration for experiment design. Different OpenSim grids will allow region administrators different levels of control in this regard, and this will oftentimes have important implications for what sorts of microeconomic systems can be implemented.[12]

- **Administrative Tools and Support:** Rather than establish a region in a pre-existing grid, it is possible to set up ones own OpenSim server or pay to have one hosted for you. While this will generally give the experimenter maximal control over the grid, the setup costs involved in doing so can be substantial. Different grids will provide different sorts of administrative tools for region management, and will also provide different levels of support if anything malfunctions. One relevant fact about OpenSim is that it is still in a beta version that experiences many bugs, and bugs can be costly if they slow down the design of experiments or occur during the experimental sessions themselves.

- **Reproducibility:** As seen with the Example Commons, one important feature of OpenSim is that it supports the ability for experimenters to export and import entire

[12] In particular, in OpenSim there are many OSSL functions that can only be used if specifically enabled by the grid's administrators. Many of these functions are fairly useful for experimental design.

regions as singular .oar files containing all the assets and code needed to run the given experiment. This allows for experiments to easily be moved from one grid to another and shared between experimental teams. From a scientific perspective, this is an extremely valuable feature; however, one must take care as not all OpenSim hosts allow for the unfettered exportation and important of .oar files.

- **Price:** Although setting up ones own OpenSim server can be done for free, this can involve having to devote time to complicated tasks relating to server administration and maintenance. For this reason among others, most users choose to pay to locate their regions on pre-existing grids. On Second Life, a single region costs $147.50 per month[13] with a special discount for educational and nonprofit organizations. However, popular OpenSim hosts tend to offer cheaper services. On Kitely, one of the more popular OpenSim-based grids, an administrator can acquire 4 regions for $19.95 per month.[14] Depending on the scope of one's experiment, an investment in multiple regions might be wise even if an experiment will only take place on a single region other regions might be useful in order to be dedicated to development. Or alternatively, if one plans to simultaneously run several instances of an experiment, using a different region for each instance may be prudent if not necessary.

At the Center for the Study of Neuroeconomics, I have used several different grids for my experiments. I began by using Second Life, which was relatively stable but was also both costly and limited in terms of its features. These drawbacks lead me towards running experiments on hosted OpenSim servers for future experiments. However, some of these grids were plagued by stability issues that lead to problems for the experimental sessions. The guidelines provided in this section can hopefully provide researchers with a set of criteria to keep in mind when choosing a grid for their particular experiments.

[13]http://specialorders.secondlife.com/
[14]https://www.kitely.com/services

2.4.2 Choosing an OpenSim Viewer

Viewers are the client software packages that are used to communicate with OpenSim servers and represent virtual worlds to subjects. As an experimentalist, using a viewer is necessary to both implement your experiment and will likely be a part of administering it as well.[15] Additionally, subjects will need viewers to interact with the microeconomic system that has been designed. An updated list of available OpenSim and Second Life viewers can be found at http://wiki.secondlife.com/wiki/Third_Party_Viewer_Directory.

Many of these viewers offer experimentalists useful tools that can be used to control avatars from the client side. As an example, consider an experiment consisting of some number of rounds where each round requires that each subject avatar start in a certain location. In OSSL this can be achieved by teleporting the subject avatars to a desired location via a script. However, if the experiment is running on Second Life, there is no equivalent function that forces an avatar to teleport. However, many viewers contain a set of features collectively called Restrained Love, which allow scripts to directly control avatar behaviors through the viewer such as executing a client-side request to teleport as if the user controlling the avatar had executed it manually via the client. Furthermore, Restrained Love can be used to disable viewer functionality that experimenters wish to prevent subjects from being able to use, such as the ability for avatars to engage in private instant messages with one another. Because of the existence of these controls and their ease of use, Restrained Love has been an integral part of maintaining experimenter control in these experiments.

One additional point of note is that many third-party viewers are either open source or customizable. This means that further hacks could be performed if necessary to build in or remove functionalities as desired. For example, even with Restrained Love activated in experiments, subjects will still generally have access to a variety of menus in the viewer's interface and can start selecting options that the experimenters may prefer they don't. Under these circumstances, editing the viewer itself to remove these menus may be desirable.

[15] The Example Commons was designed by an inworld avatar using a viewer. Furthermore, an administrator avatar is needed in world to start and stop experimental sessions.

2.4.3 OpenSim and the Web

As mentioned earlier, it is possible for OpenSim to communicate with remote servers using standard HTTP protocols, essentially allowing for each script to have its own API. This provides for a lot of flexibility in design that my experiments have only barely explored. For example, it is possible to have subjects in OpenSim interact with subjects using other experimental platforms such as Amazon's Mechanical Turk or zTree. It is possible to pull in all sorts of web data to be presented in an experiment, if desired.

Additionally, the use of script APIs can be used to circumvent some of the clunkiness of OSSL itself as well as having to manage complex object communications within the virtual world. Given this, one may wonder why I haven't recommended simply having most of the processing for an experiment done via a web application that replaces much of the communication structure depicted in Figure 2.12. For example, in the Example Commons, why have I implemented a master object instead of a remote script programmed in a language of my choice that communicates via HTTP requests with the berry bushes?

There are two answers to this question: The first one is simply to observe that even were things this simple, this would require familiarity with several additional software tools and languages. Learning OSSL is already fairly difficult, and I do not wish to impose additional burdens on the reader by assuming any experience with the myriad of topics involved with web development.

The second is that, in my experience, using APIs within scripts is not quite as powerful as the capability may first appear. In particular, Second Life and many OpenSim grids throttle the rate at which HTTP requests can be made and oftentimes the latency of communication will be substantial. This may make it infeasible for, say, many berry bushes to quickly communications about their status with a centralized server in a timely fashion. If an experimenter wishes to implement an experiment that involves extensive use of OSSLs API functionality, I would advise them to take these limitations into consideration when exploring possible communication structures. Researchers who do not feel comfortable with these tools should not feel discouraged if they cannot be used.

2.5 Administering an OpenSim Experiment

Once you've selected your software tools and implemented your microeconomic system, the final step is of course to actually execute it with subjects. This section is dedicated to outlining advice related to acting as an administrator for an OpenSim experiment and maintaining experimental control during a session. This advice will be primarily focused on laboratory experiments where subjects are using a virtual world, but much of it will also be useful for other experimental setups.

2.5.1 Experiment Preparation

OpenSim viewers are more hardware-intensive than normal web-based experiments or zTree experiments, and it's important to make sure the hardware being used by subjects is capable of running OpenSim in a reasonable fashion. With that said, most desktop computers less than a decade old should be able to run the software adequately. The graphics settings on OpenSim can be adjusted to improve performance if necessary, and should be standardized across all the computers being used to ensure that no subjects have innate advantages in the virtual world arising from using computers with better hardware or software settings.

It is crucial to extensively test your experiment before proceeding with data collection. This testing should not only concern whether the core functionality of your experiment appears to be working properly, but also how it handles several users interacting in your environment at once complicated interactions may cause lag or server congestion that was not previous anticipated and must be addressed. For example, in the Example Commons experiment the number of berries on a berry bush was represented through a single texture that would update if the bush's population changed (see Figure 2.6.) Naturally, part of testing that this is working properly would be to test one berry bush, but if we had a sufficiently large number of berry bushes and many of them were being simultaneously updated this would produce a noticeable amount of slowdown. The only way to properly prepare for these possible issues is to pilot your experiment under conditions as similar as possible to what your real sessions will look like. If you're going to be running multiple groups at once during a session, see if your grid can handle multiple groups all running at once.

During these pilot sessions, it is important to get feedback from your subjects on whether the environment was working as described in their instructions. These environments are often very complex and it's possible that subjects have discovered unanticipated strategies or interactions that you may wish to address. For example, in one implementation of the Berry Island experiment I ran multiple groups at once, each having their own region to interact on. However, some subjects managed to find a way to get over the wall separating the different regions and managed to enter the experimental environments of the other groups, which contaminated the data of all the groups involved. The fact that this crossing had occurred was immediately apparent in the raw data and thus detected by the experimenter, but it wasn't until this matter was investigated with the subject's help that the exact cause of the problem was identified and fixed.

2.5.2 Subject Training and Comprehension

One common concern that Ive heard raised about using virtual worlds in a laboratory setting concerns subject training how much time does it take to impart a basic familiarity with controlling an OpenSim avatar in general, and how difficult is it for subjects to understand particular experimental environments?

To the first question, in my experience most laboratory subjects drawn from a college population are able to immediately grasp the basics of controlling an avatar moving an avatar is simply a matter of using the keypad, and interacting with experimental objects involves clicking on either the HUD, the environment, or typing messages into a chat window. Usually covering these concepts takes a couple minutes of explanation at most.

The question of training on particular environments, however, is obviously more complicated and the ease of doing so will depend on many of the particulars of the given environment. Ideally, a well-contextualized microeconomic system should have many features that may be more intuitive in a virtual world than elsewhere for example, in my experience its easy to explain to subjects the importance of objects like the berry bushes in the Example Commons and how to interact with them. Explaining the particulars of the logistic growth function may pose a challenge, but this is not a challenge that stems from the use of the virtual world itself.

One point to keep in mind is that if subjects will be interacting in the virtual environment in real time, then it is important to not have an interface that is cumbersome or unintuitive. Not only will this tend to necessitate additional training time, but clunky interfaces could serve to undermine parallelism by drawing the attention of subjects away from the events occurring in the microeconomic system. For example, in my early experiments performing actions such as transferring resources would tend to involve having to interact with several menus first the subject would have to choose which resource to transfer, and then choose how much to transfer, and then who the recipient of the transfer was. This meant that subjects would often have to stop for a considerable amount of time to execute what should be a fairly simple action. Future experiments involving resource transfer streamlined this process in various ways.

An important related concern is that complex interfaces may lead to confounds in experimental outcomes based on differentials in familiarity with the software. A subject who is able to quickly navigate the aforementioned complex transfer system will be able to execute more transfers, which may make him or her a more-desirable trade partner within the experiment. These sorts of concerns are certainly not confined to virtual world experiments, but are more likely to manifest under designs that feature continuous-time decisionmaking with complicated interfaces.

2.6 Conclusion

Virtual worlds offer a great deal of potential for the implementation of designs that explore new and exciting research questions in controllable, reproducible environments. Their naturalistic contexts and large native populations make them attractive as settings for a variety of types of field experiments. Despite this potential, however, experimental economists have only scratched the surface of what can be done with virtual world experiments.

To a large extent, this avoidance is understandable implementing microeconomic systems in virtual worlds often requires experience with unfamiliar software tools. As such, a major goal of this paper is not merely to convince the reader of the fundamental value of these experiments but also to provide some constructive direction on picking up one of the most popular virtual world platforms that has been used by experimentalists, OpenSim. I believe that OpenSim and open platforms like it will continue to have a growing impact on the experimental literature by providing a valuable bridge between the laboratory and the field.

Chapter 3: The Determinants of Territorial Property Rights in a Spatial Commons Experiment

This study explores costly specialization in a heterogenous common pool resource environment, allowing for the investigation of the emergence of two different types of property right systems as discussed in Demsetz (1967). Demsetz refers to these two systems as communal and private property rights. In a communal property rights system, resources are owned by the group as a whole who must then agree to monitor and enforce the rules which govern the commons. These rules by necessity must include detailed operating procedures or usage rules. In a private property rights system, the group must agree to and enforce ownership rights to individuals, allowing them to individually choose operating procedures. The group must then adjudicate disagreements caused by any externalities or enroachments that result.

Do either of these systems work in practice? In her field research on common pool resource (CPR) management, Ostrom and her colleagues developed the Institutional Analysis and Development (IAD) framework to help identify features shared by successful CPR regimes (Ostrom, 1990, 2000). This framework resulted in the discovery of design principles such as clear boundaries on usage rights, monitoring to detect rule violations, and the use of graduated sanctions on rule violaters. Furthermore, this research uncovered a variety of different interactions between members of the group as they tailored these design principles to changing environments.

The analysis of these design principles has generated a large body of experimental research as social scientists have turned to the laboratory to study which principles are necessary and/or sufficient for the management of CPRs or other closely-related social dilemmas. When the environment is easy to understand, often only one of the principles is sufficient.

Classic examples of this include communication (Bochet, Page and Putterman, 2006; Brosig, Weimann and Ockenfels, 2003; Isaac and Walker, 1988) and punishment (Carpenter, 2007; Denant-Boemont, Masclet and Noussair, 2007; Nikiforakis and Normann, 2008). However, when Janssen and Ostrom (2008) studied CPR management in a more natural spatiotemporal environment, where the CPR stock is replenished according to a biologically inspired logistic population growth model, they discovered that these design principles are no longer individually sufficient. Their experiment showed that neither punishment nor communication on its own was effective in overcoming the tragedy of the commons; rather, subjects had to have access to both of these tools to avoid overharvesting.

In laboratory experiments designed to implement more field-like settings, this result raises the following question: When do treatment conditions in simple, controlled laboratory environments and the models tested therein explain how people make decisions in naturally occurring environments? This is the parallelism condition suggested by Smith (1982) as one of the sufficient conditions for an economic experiment. This growing interest in parallelism has lead economists to run field experiments to further refine and test theories (Harrison and List, 2004). While extremely valuable, field experiments can be very costly and difficult to replicate.

In this paper an experiment was designed that uses a virtual world infrastructure. Subjects interact with the environment and each other through individual avatars that represent their virtual self to others in the experiment. The virtual world laboratory allows an experimenter to closely duplicate natural mechanisms for social exchange and create environments that induce natural spatiotemporal reasoning. This allowed for the construction of a commons environment more like those studied in the field which lead to the IAD framework and the discovery of Ostrom's design principles for successful commons management. The primary advantage of this approach is that it provides the ability to enforce experimenter controls, allowing for replication and follow-up investigation at reasonable costs. The disadvantage is that this approach causes more variation and noise in the decisions of subjects. The open question is do virtual worlds come close enough to replicating field conditions while maintaining experimental control?

In this experiment I chose to implement a design which modified the Janssen and Ostrom (2008) CPR environment to allow for a heterogenous CPR resource (called berries) differentiated by color that can be harvested by subjects. Berries grow on berry bushes which replenish independently according to a stochastic growth function that depends on the number of berries on that bush. Groups of subjects are given a fixed amount of time to harvest berries. I hypothesize that myopic self-interest will lead to a tragedy of the commons. To allow subjects to manage this problem, they were given the same social tools found to increase success in the Janssen and Ostrom experiment, namely the abilities to communicate with and punish each other. To study the emergence of different property right systems, mutual monitoring was made imperfect in order to increase the difficulty in enforcing any agreed-upon rules of harvesting or rights to use resources.

This study investigates in particular determinants of whether groups establish private or communal systems of property rights. One common system of private property rights that can be established in this environment is that of territoriality, wherein subjects engage in resource harvesting within individual non-overlapping spatial regions. Gintis (2007) points to the commonality of territorial behavior in various animal species to argue that humans have an ingrained predisposition towards establishing territorial property rights, while DeScioli and Wilson (2011) show that territorial ownership conventions can emerge in laboratory experiments even in the absence of robust natural-language communication. The lessons of the literature on human territoriality as a CPR management tool are not confined to the foraging settings similar to the ones implemented in this experiment, but have been applied to topics such as how usage rights over resources are established within firms (Avey et al., 2009; Brown, Lawrence and Robinson, 2005) or in collaborative spaces such as Wikipedia (Thom-Santelli, Cosley and Gay, 2009).

Two different treatments are studied that I hypothesize will modulate both the overall quality of CPR management and also the propensity of groups to establish territorial property rights. The chosen layout treatment varies the layout of berry bushes in the virtual berry field. The cost of specialization treatment controls how costly it is for subjects to specialize, ie. receive an enhanced payoff from berries of particular colors. In the layout treatment, two berry fields are considered. One field, which I call "mixed", has a less differentiated boundary between different types of berry bushes. The second field, which I call "separated", has a more-differentiated boundary between types of bushes, making it easier, I hypothesize, to monitor and enforce property rights. This treatment variable is inspired by the IAD's emphasis on clearly-defined boundaries and ease of monitoring as important factors in successful CPR management regimes. As it turns out I observed no statistical effect of this treatment on groups' ability to manage the commons as measured by total group harvest within a period or on territoriality.

The specialization treatment implemented either no cost specialization or costly specialization. In the no cost specialization condition, subjects could costlessly switch their specialization when harvesting berry bushes to obtain an enhanced payoff for harvesting that bush. Under the costly specialization condition, subjects could only commit once within a round to receive an enhanced payoff from a single color variety of berry bush. Because this enhanced payoff could be realized without reducing the payoff received from other berry types, choosing a specialization represented a weakly-dominant strategy for individuals. When subjects made a specialization decision in either condition, their avatar's shirt color changed to the color of the berry bush, allowing other subjects whose avatars were nearby in the spatial environment to observe this choice. Specialization has been recognized as a crucial feature of organized societies since the contributions of Smith (1776). In more recent times, biologists and ecologists have recognized the role that specialization plays in the formation of ecological niches as a response to both interspecific and intraspecific competition over common pool resources (Futuyma and Moreno, 1988; Pyke, 1984). It is hypothesized that this signal to other subjects would have more impact in the costly specialization condition since it signals a commitment by a subject to a resource type that would foster the development of a private property system to manage the CPR. In Thom-Santelli, Cosley and Gay (2009), it is shown that Wikipedia editors establish property rights over different articles based on relative expertise. The broader question is whether the presence of this sort of specialization in a commons environment can serve as a basis for establishing private property rights regimes that are conducive to good commons management. In this experiment it is shown that costly specialization significantly improves the management of the commons, and that this improved management coincides with the establishment of stronger territorial institutions. This study demonstrates that small groups can quickly move to private ownership of land when their territorial instincts are activated by costly specialization.

3.1 Experiment and Hypotheses

3.1.1 Harvesting Task

This experiment is built around a 15-minute harvesting task. In this task, a group of 4 subjects decides how to harvest a field of 48 berry bushes subdivided into five color types (red, green, blue, yellow, and silver). Each bush can support from 0 - 9 berries which grow according to the following rule: If a bush has m berries at time t, then the probability of that bush having n berries at time $t+1$, denoted H_{mn}, is calculated using the first-order Markov transition matrix H defined as follows:

$$H_{mn} = \begin{cases} 0 & , m > n \\ \binom{10-m}{n-m}(1-\min(\phi(m-1),1))^{10-n}(\min(\phi(m-1),1))^{n-m} & , m \leq n \end{cases} \tag{3.1}$$

The real-valued parameter $0 \leq \phi \leq 1$ monotonically adjusts this overall rate of growth. For values of $\phi < .1$, this expected growth rate is maximized when there are exactly four or five active berries, as shown in Figure 3.1.[1] This quadratic growth rate illustrates the

[1] A discussion of the implications of higher values of ϕ is contained in Appendix A.

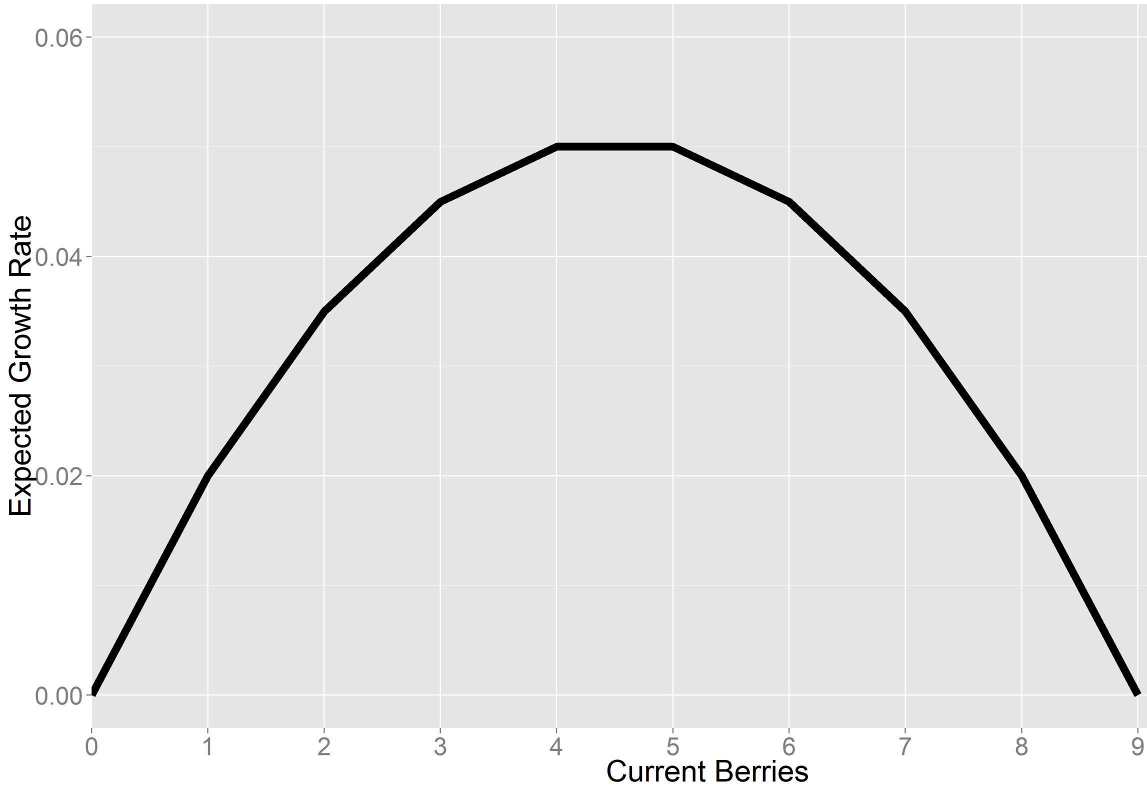

Figure 3.1: Expected berry bush growth per period as a function of current berries when the growth parameter $\phi = .0025$.

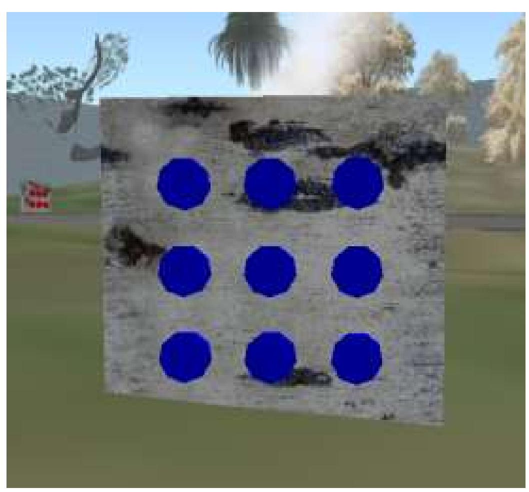

Figure 3.2: Blue berry bush with 9 active berries

necessity of commons management in order to prevent overharvesting, resulting in lower total resource yields. A typical blue berry bush with 9 berries is shown in Figure 3.2. In this experiment, a subject can pick berries from the bush by navigating their avatar to within 5 meters of the bush and then clicking on the berry to be picked. Subjects' berries are converted into money at the end of the experimental session, giving them an incentive to manage the common pool berry resource in order to maximize income.

In the experiment, subjects collect berries in one of the two spatial layouts depicted in Figures 3.3 and 3.4. This figure shows both the mixed and separated layouts used in this experiment, with the colored squares representing the locations of berry bushes of the corresponding color. In both treatments there are 24 silver berry bushes and 6 red, 6 green, 6 blue, and 6 yellow (RGBY) bushes. In both layouts the berry bushes were grouped together by color to facilitate the formation of territorial boundaries. However, in the separated layout the RGBY bushes were clustered more tightly together while the silver berry bushes act as a perimeter around the RGBY bushes.

I first establish some simple benchmarks that assume subjects in the harvesting task possess perfect information and can harvest the field to any extent in each period. For the purposes of the simplified harvesting task the state of this field at any time t is represented through the vector $s^t = (s^t_1, s^t_2, ..., s^t_{48})$ where s^t_k is the number of berries on bush k at time t. If $\mathcal{S}$ is the set of all possible states then $\mathcal{S}$ has 10^{48} elements. It is assumed that in each period once a subject i observes s^t then i makes a harvesting decision $b^{it} \leq s^t$ where b^{it}_k is the number of berries picked from bush k.

Under no cost specialization, subjects can choose to specialize over any combination of colors excluding silver. If subject i is specialized in a given color then each berry of that color collected is worth $r\rho$ dollars, where $r > 1$ represents the monetary premium for being specialized in a color type. If subject i is not specialized then each berry collected is worth ρ dollars. A subject trying to maximize payoffs in this game will always choose to be specialized in red, green, blue, and yellow bushes in this treatment. This allows for the specification of a static vector δ^i such that $\delta^i_k = 1$ indicates subject i chose to be specialized

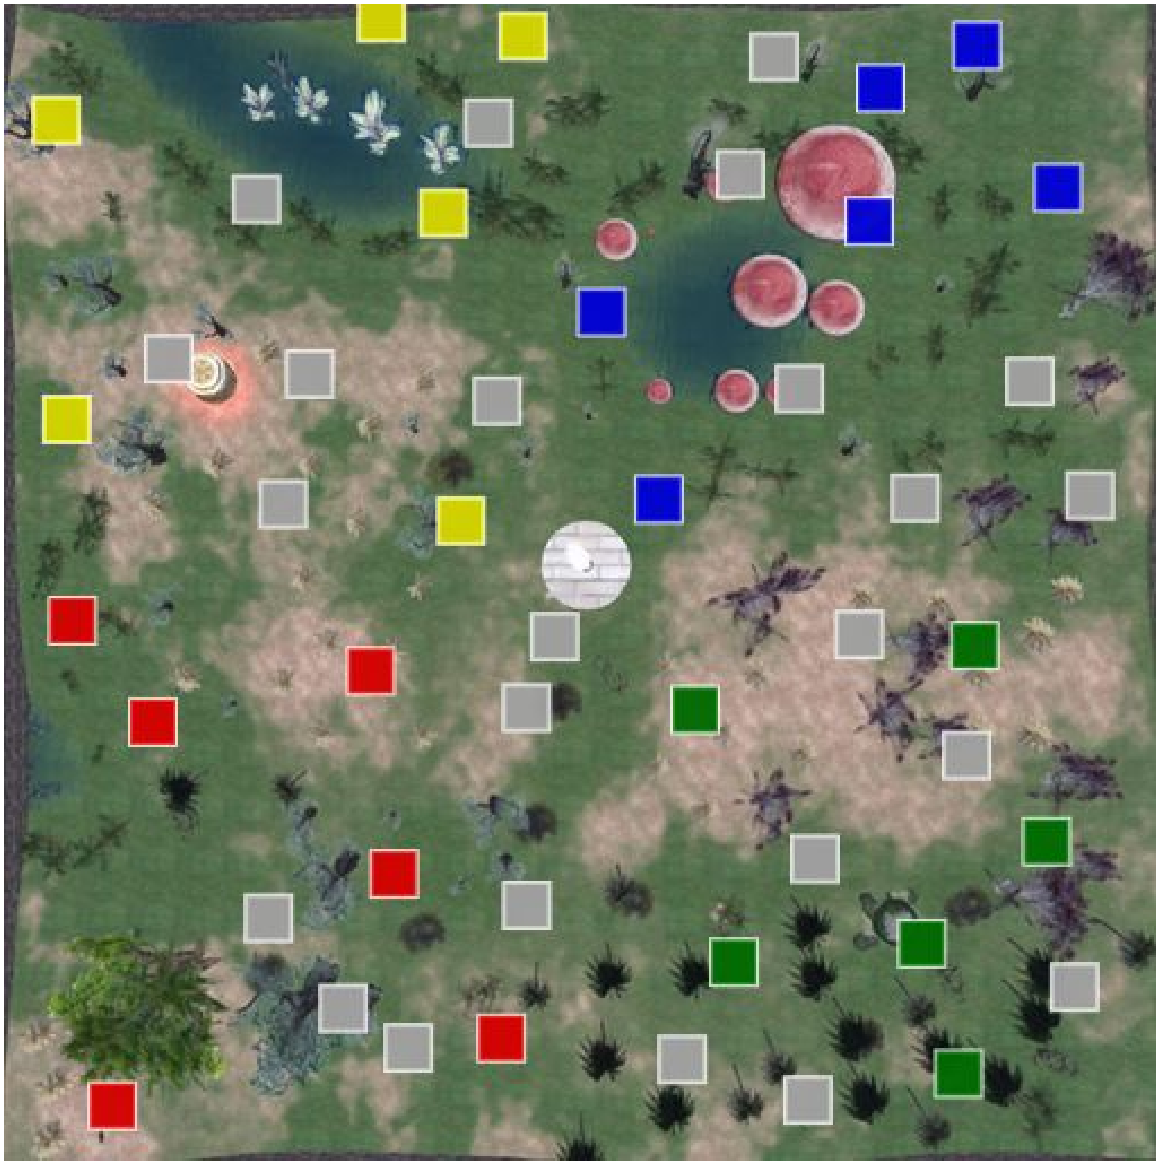

Figure 3.3: Mixed layout

in the color type of bush k and $\delta_k^i = 0$ otherwise. Note that δ_k^i is constrained to be 0 for all silver bushes.

Under costly specialization subjects may only choose to specialize once during the round and only in one color other than silver. The set of possible specializations that can be chosen under this treatment is denoted as $\Delta = \{\delta(k) \text{ where } k \in \{red, blue, green, yellow, none\}\}$

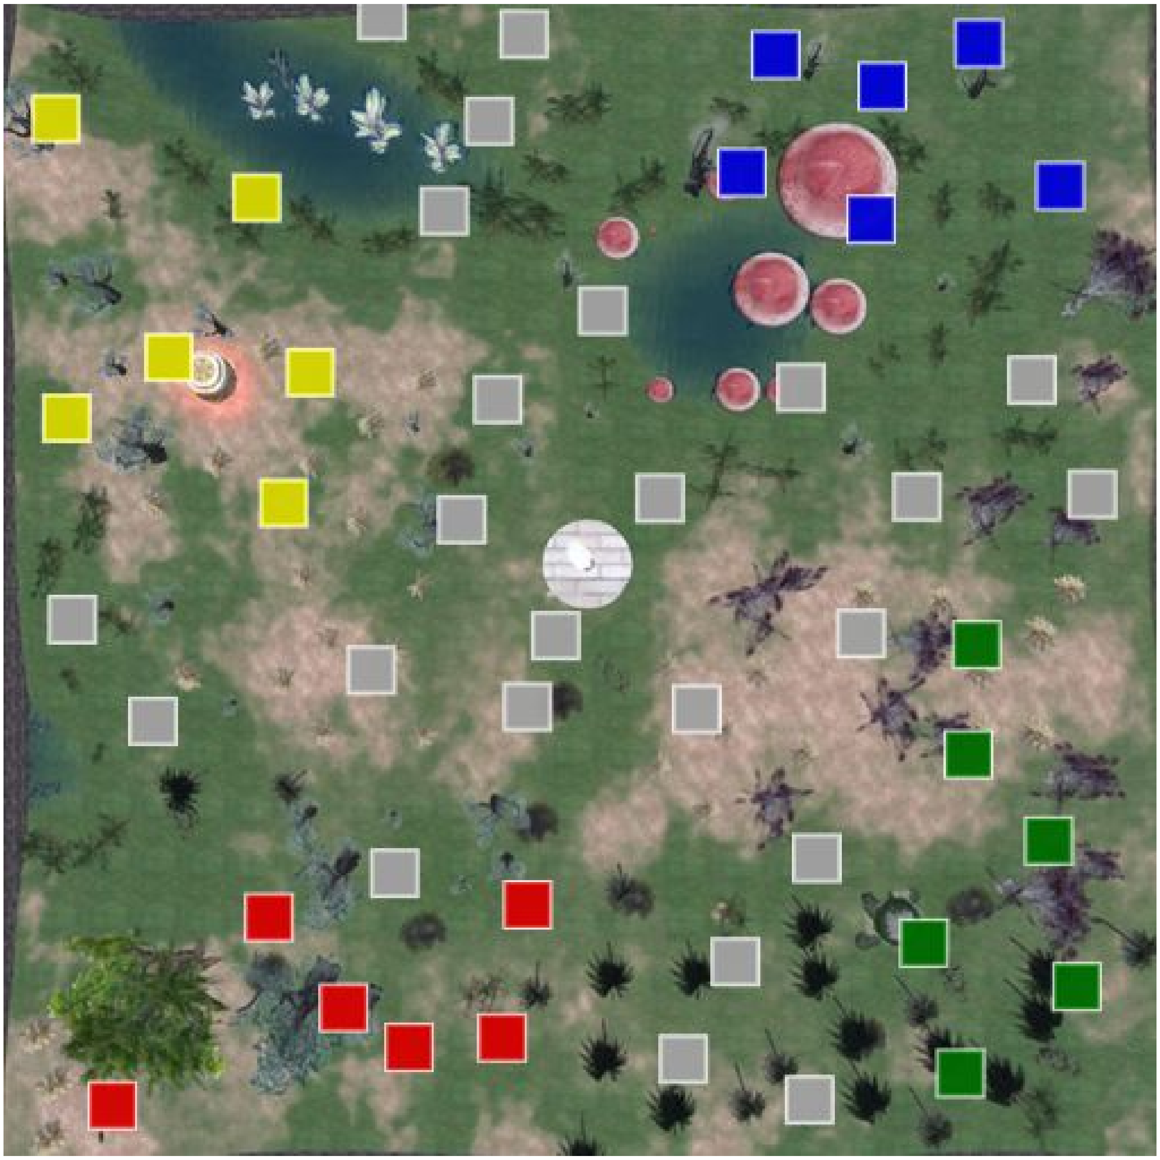

Figure 3.4: Separated layout

such that $\delta(none)$ means that $\delta_j = 0$ for all j and $\delta(red)$ means that $\delta_j = 1$ only if j represents a red bush, etc. It should be clear that a profit maximizing subject will pick $\delta(k)$ where $k \in \{red, blue, green, yellow\}$ as soon as possible.

The payoff for a subject (dropping the subject index subscript for now) in a single period t as a function of a harvesting decision b^t and a specialization δ can now be expressed as:

$$\Pi(b^t, \delta) = \rho \sum_{k=1}^{48} [r\delta_k b_k^t + (1 - \delta_k) b_k^t] \tag{3.2}$$

I begin this analysis by asking what policy will maximize this payoff. To answer this question I set up the associated finite time, discrete state, Markov decision model for a single decision maker which we will call the planner. Using dynamic programing, it can be demonstrated that the optimal policy for the planner is as follows: At every point in time t, the planner looks over the field to observe the state of the field s^t, which consists of the number of berries on each bush. Using Bellman's principle of optimality (Bellman, 1957) the planner will solve the problem

$$V^t(s^t) = \max_{b^t \leq s^t} \{\Pi(b^t, \delta) + \sum_{s^{t+1} \in S} P(s^{t+1}|s^t, b^t) V^{t+1}(s^{t+1})\},$$

$$s \in S, t \in \{1, 2, ..., T-1\} \tag{3.3}$$

As shown in Appendix A, a solution to this problem involves determining an optimal berry cutoff c^* and time period τ^* such that,

$$\forall k \in B,\ b_k^t = \begin{cases} 0 & , s_k^t \leq c^* \text{and } t < \tau^* \\ s_k^t - c^* & , s_k^t > c^* \text{and } t < \tau^* \\ s_k^t & , t \geq \tau^* \end{cases} \tag{3.4}$$

For the chosen value of the growth rate parameter $\phi = .0025$ the optimal planner chooses $c^* = 4$ and $\tau^* = T$. In other words, the planner does not clear the field until the terminal time period T, and until that point they will choose to harvest berry bushes down to 4 berries whenever they grow above this c^*. $\tau^* = T$ is intuitive since the total number of berries will always grow from one period to the next, and thus clearing the berry bushes at a time $\tau < T$ will simply forgo the opportunity to harvest the berries that may have grown after τ. The cutoff $c^* = 4$ reflects choosing a value that optimizes marginal berry growth in each time period.

A rule that chooses a cutoff c and a clearing time τ will be called a management rule of the form $g(c, \tau)$. Now let's imagine the non-cooperative commons game where, unlike the planner's problem, each subject in each groups will try to maximize their own income. Each subject i can be characterized as having their own management rule $g(c^i, \tau^i)$. By backwards induction wa best response by any subject i to any subject j's clearing the field at time τ^j is to clear the field at time $\tau^j - 1$, resulting in a Nash equilibrium policy where the field is cleared immediately at $\tau^k = 1$ for all k. This results in fewer harvested berries for subjects than if they all followed $g(c^*, \tau^*)$. Thus, the planner policy and the Nash equilibrium policy determine the upper and lower bounds on what can be achieved through active management of the CPR.

As an example, let's imagine the case where $\phi = .0025$, $T = 180$, and each berry bush is initially randomized so that $s_k^0 \in \{1, 3\}\ \forall k \in \{1, 2, ..., 48\}$. Under these assumptions a group following the Nash rule $g(c, 1)$ will harvest an expected 96 berries, while a group following the planner rules $g(4, T)$ or $g(5, T)$ can expect to harvest an average of approximately 486 berries.[2] This differential represents the tragedy of the commons in the harvesting task. These were the parameters used in the actual experiment faced by subjects.

Given the discussion above it can be seen that managing the CPR requires some mechanism for governing subjects so that they follow a management rule $g(c', \tau')$, where $c' > 1$ and $\tau' > 1$.[3] At this point one can distinuish between Demsetz's communal property rights and private property rights systems. If subjects decide to treat the berry bushes as communal property, they have to mutually negotiate the rule $g(c', \tau')$ and try to enforce it for as long

[2]This approximation is derived by simulating the growth of the environment under subjects employing a planner strategy. 10000 trials were used to derive a mean harvest of 486.03 berries with a standard deviation of 22.05 berries.

[3]It should be clear that c' values closer to 4 and τ' values closer to T produce higher income, but for various right negotiation and enforcement reasons groups may pick (c', τ') values in intermediate regions.

as possible. If subjects agree to use a private property rights system, they only have to agree on who has exclusive rights to which bushes, but they can leave the choice of management strategy $g(c^i, \tau^i)$ up to each individual since there are no harvesting externalities between bushes.

3.1.2 Hypotheses

The analysis from the previous section indicates that if subjects manage the commons in a near-optimal manner then they will produce more berries than if they fall victim to the tragedy of the commons and clear the bushes as soon as possible. Within this framework I have three viable hypotheses. My first hypothesis, which I call the social interaction hypothesis, predicts that if subjects are given the social tools necessary to implement Ostrom's design principles, such as punishment and natural language communication found to be effective in past experiments, then they will be able to effectively manage the CPR. If this hypothesis is accepted it would lead to the conclusion that our current understanding of effective CPR management, as studied in Janssen et al. (2010), is sufficient at least for environments similar to those of this experiment. My second hypothesis, which I call the costly signaling hypothesis, predicts that management of CPRs will improve in environments where there is a cost to specialization. If this hypothesis were accepted it would lead to the conclusion that further research is warranted on the role of costly specialization in defining and maintaining different types of property rights that allow better CPR management. My third hypothesis, which I call the clear boundary hypothesis, predicts that if different berry colors are more-clearly delineated spatially and easily monitored better CPR management will be observed. If this hypothesis were accepted it would lead to the conclusion that further research is warranted on how clearly defined boundaries are created and how such boundaries interact with the ability to monitor participants in CPRs.

I will state these hypotheses in terms of total berry production $R^{\alpha\beta}$ where α, the field layout treatment variable, is equal to m in the mixed field layout condition and equal to s in the separated field layout condition, and β, the specialization treatment variable, is equal to c for in the costly specialization condition and n in the no cost specialization condition. With respect to total berry production I now define the hypotheses:

Hypothesis 1.

$H_0^1 : R^{mc} = R^{mn} = R^{sc} = R^{mn}$

$H_1^1 : R^{mc} + R^{sc} > R^{mn} + R^{sn}$

$H_2^1 : R^{sc} + R^{sn} > R^{mc} + R^{mn}$

If the social interaction hypothesis is correct the null hypothesis H_0^1 will be accepted. However, if the costly signaling hypothesis is correct, H_0^1 will be rejected in favor of H_1^1. Finally, if the clear boundary hypothesis is correct, H_0^1 will be rejected in favor of H_2^1.

I also consider one mechanism for CPR management which I call territoriality. Territoriality asserts that when subjects maintain uncontested territories, they will be able to better manage their CPRs. Therefore, a measure of the size of a group's uncontested territory, M, will be expected to grow with better CPR management. This leads to the following hypotheses:

Hypothesis 2.

$H_0^2 : M^{mc} = M^{mn} = M^{sc} = M^{mn}$

$H_1^2 : M^{mc} + M^{sc} > M^{mn} + M^{sn}$

$H_2^2 : M^{sc} + M^{sn} > M^{mc} + M^{mn}$

If one rejects H_0^1 but accept H_0^2 then it must be concluded that, while the social interaction hypothesis is insufficient, it is not the case that subjects achieve better CPR management through the establishment of uncontested systems of territorial private property rights.

3.2 Experiment Procedures and Methods

Using virtual world platforms to run experiments is a relatively new undertaking. OpenSim was chosen as a platform for this experiment because it provides several built-in development

Figure 3.5: Subject interface

features that are desirable for the construction of naturalistic online experiments.[4] As shown in the experiment screenshot in Figure 3.5, OpenSim environments are intrinsically spatial, are avatar-based, and allow for the easy implementation of subject chat and spatial monitoring.[5] Unlike the simplified harvesting task, subjects will have imperfect knowledge of s^i constrained by their ability to monitor the spatial setting, and will have to spend time traveling between different berry bushes to engage in harvesting.

Interfaces of this type are becoming common in experiments concerning endogenous institutions (DeScioli and Wilson, 2011; Janssen et al., 2010; Wilson et al., 2012), as they provide a natural context within which institutional formation can be grounded. Several studies have shown that the naturalistic framing of economic dilemmas can have important effects on subject behavior, often mediated through the modulation of social distance.[6] We

[4]In order to run a virtual world experiment, the experimenter can build or buy virtual world objects that are then scripted to provide specific interaction behaviors with subject avatars. While there is a relatively steep learning curve in learning how to build and script objects, an experienced experimenter can build a virtual world experiment in the same time frame it takes to program other economics experiments.

[5]While this makes social interaction more realistic, as of now social interaction in OpenSim does not include facial expressions and body language that may help subjects infer each other's intentions. Developing these capacities is currently an active research program in OpenSim.

[6]See Al-Ubaydli et al. (2013); Charness, Haruvy and Sonsino (2007); Hoffman, McCabe and Smith (1996);

are only beginning to learn how to take full advantage of virtual world technologies while simultaneously allowing for straightforward replication. I have broken down our discussion of these considerations into two sub-sections. First, I discuss the decisions that were made in building this virtual world experiment. Second, I discuss the decisions that were made that allowed for the maintainence of experimenter control in the experiment.

3.2.1 Mechanisms for Social Interactions

Establishing property rights requires social interaction. Many mechanisms for facilitating social interaction, in order to improve cooperation, have been studied in the laboratory. This experiment includes many of these common interaction mechanisms including natural language conversations, common or focal places to meet, costly punishment, identity, and transfers or gifts. These mechanisms are introduced in a way that I hoped would seem natural to subjects as explained below.

This experiment allowed subjects to use social mechanisms that are commonly available to human social groups. Subjects engaged in conversation through typed natural language messages. Such conversations were localized and had to take place between avatars within a twenty meter radius of each other. Furthermore, naturally occurring groups have meeting areas that facilitate group conversations. In this experiment a meeting area was built to serve this purpose. This meeting area was a twenty meter-wide circular platform with a central campfire that was placed in the middle of each field (see Figure 3.5). To incentivize subjects to use the meeting area, they earned .048 dollars per minute whenever they stood on the platform. While this amount was small, it helped to establish the meeting area as a natural place to converse as a group. It further serves as a place to mutually monitor other subjects; i.e., if someone is at the center then they are not poaching berries, and when someone leaves the center it is easy to track their path to see if they are moving towards someone else's bushes.

Kimbrough and Wilson (2011) for general examples. Additionally, Atlas and Putterman (2011); Fiedler, Haruvy and Li (2011) specifically examine framing and social distance issues in environments based on Second Life, a platform closely-related to OpenSim.

Additionally, subjects could also punish one another. In the experiment, subjects were given 200 punishment tokens at the beginning of each round. Each token cost .004 dollars to use and resulted in a .008 dollar loss to the subject being punished, meaning that 1.60 dollars in earnings could be deducted from other subjects by any subject in each round. A further limitation on the use of punishment tokens is that a subject's avatar had to be within a twenty meter radius of another subject's avatar in order to punish that subject. This raises the difficulty (or cost) of punishment since a subject who is poaching can simply run away if they see the 'owner' of the bush approaching. The owner would then have to decide whether or not to spend precious time chasing down the offender in order to punish them. Finally, it should be noted that the threat of punishment may keep offending subjects away from the commons area and out of the group conversation.

When subjects harvested berries they acquired a number of seeds of the corresponding color. The amount of seeds they acquired was dependent upon whether or not they were specialized in the given color. Each seed was worth .016 dollars. The parameter values of the baseline payoff per berry was set to $\rho = .016$ and the specialization premium was set to $r = 2.5$, meaning that subjects who harvested a berry of a color they weren't specialized in would receive a quantity of seeds worth .016 dollars while harvesting a berry of a color they were specialized in would would yield .04 dollars.

Subjects were also allowed to transfer seeds to one another. Such transfers can be used to reward or reciprocate good behavior, to attempt to bribe an avatar to behave in a certain way, or as a payment for bad behavior. Desmet, Cremer and Dijk (2011) shows that monetary transfers can serve as an important facilitator of trust repair after a perceived norm violation. In the experiment, if one is caught poaching it might make sense to hand the berries over to the owner. Such behavior would reduce costly punishment and the risk of being informally ostracized from the meeting area. As an example of resource transfers facilitating trust repair, in his discussion of private property rights among the Montagnes Indians Leacock (1954) mentions that "a starving Indian could kill and eat another's beaver if he left the fur and tail." Note that the fur and tail were the more valuable parts due to their trading value.

In the naturally occurring world, people have features (both physical and legal) which allow them to be identified, and it is frequently the case that offenders try to hide their identity to escape punishment. In this experiment were not allowed to change their avatar's identity. All avatars looked identical in appearance and also when doing the same motion; i.e., standing, walking, running, gesturing, and jumping. However, each avatar also had a unique name which appeared in small print above the avatar's head. The avatar name could be read if another subject's avatar was within twenty meters. Furthermore, during chat, the avatar played a typing animation when typing out their message and all avatars in range of the message saw the name of the avatar who sent the message. Thus within the range of twenty meters an avatar could be identified. Beyond that range, avatars could be tracked only through the relatively continuous monitoring of their path from when they were within range or by process of elimination; e.g. if three avatars are in the meeting area then the fourth avatar who isn't there can be identified. Issues of identity will therefore be important in the experiment to identify rule violators and punish them. Even if a rule violator is identified they must be found within the spatial environment in order for this punishment to occur, which represents an additional difficulty for potential punishers.

3.3 Data Analysis

Data was collected from eight twelve-person sessions that took place in the Center for the Study of Neuroeconomics laboratory at George Mason University during the Spring 2011 and Fall 2011 academic semesters. Each session lasted for three hours, starting with an hour of software training[7] followed by eight 15-minute rounds of the experimental procedure. In each session the 12 subjects were first divided into three groups of four. Subjects interacted with the same group in rounds 1-4. Groups were reshuffled after the fourth round and subjects interacted with their new groups in rounds 5-8. I call this first set of four rounds the first stage of the experiment and the second set of four rounds the second stage of the

[7]Subject instructions can be found in Appendix B.

experiment. The specialization treatment condition was randomized across sessions while the field layout was randomized across stages: Each subject experienced the mixed field layout in one stage and the separated field layout in the other. The ordering of the layouts was blocked across sessions, ensuring that half the groups encountered the mixed field layout during the first stage and half of them encountered it during the second stage.

Data from the experiment is aggregated within a four-person period to provide observations of the form $Y_{t,s}^{i,j}$, where i, j are the specialization and field layout treatments variables, $t \in \{1, 2, 3, 4\}$ is the round number, and $s \in \{1, 2\}$ is the stage. To test my hypotheses, the chosen identification strategy uses the mixed effects model shown in equation 3.5 based on round-level data. The data from this experiment consists of 192 observations clustered into 48 groups, with each group providing four rounds of data. The estimators discussed later in this section are derived from equation 3.5:

$$Y_{t,s}^{i,j} = a + b^{MF} MF_i + b^{CS} CS_i + b^{\text{Order}} O_i + \sum_{j=2}^{4} b^j T_i^j + BX_i + u_{g(i)} + \varepsilon \tag{3.5}$$

Where:

- MF_i is an indicator variable that is equal to 1 under the Mixed Field layout.
- CS_i is an indicator variable that is equal to 1 under the Costly Specialization treatment condition.
- O_i is an indicator variable that is equal to 1 if the observation comes from the second stage of a session.
- T_i^j is an indicator variable that is equal to 1 when the current round is equal to j.
- X_i refers to additional covariates that may be employed in a particular test.

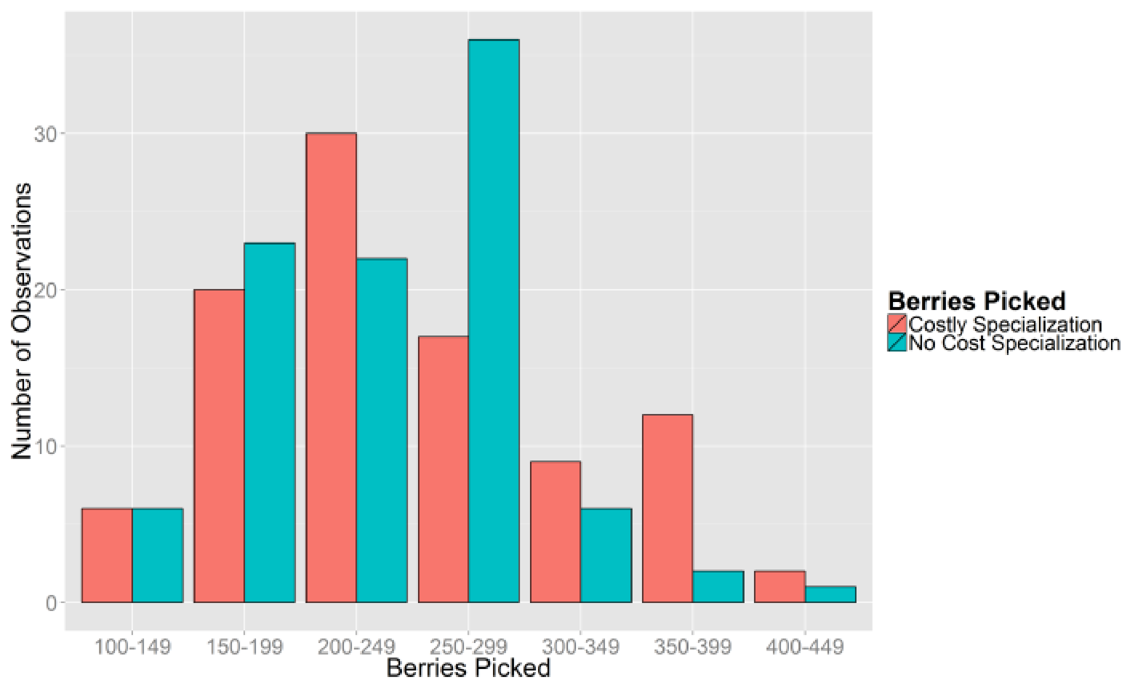

Figure 3.6: Frequency histogram of the total number of berries picked in each group, by specialization condition

- $u_{g(i)}$ represents a group fixed effect estimate for observation i. Since each of the 48 group combinations are used to generate four observations, this estimate should capture within-group correlations in the outcome variable.

3.3.1 Berry Harvest Yields

The first set of observational measures looks at the total berries harvested by the four participants in each observation, labeled $R^{i,j}_{t,s}$. In Figure 3.6, I plot the histogram of total berries harvested by specialization condition. In general there is a wide variance in overall performance across these treatment conditions. Of the 31 observations that harvested more than 300 berries, 22 of them occurred in the costly specialization condition and 19 of them occurred in the second stage of the experimental session. Although no observations saw groups achieve the average maximum of 486 derived under the simplified model, there were several observations where 400 or more berries were picked, showing that the frictions

Table 3.1: Treatment effects on berry production

	R	R_C	R_S
Intercept	226.807^{***}	108.469^{***}	118.339^{***}
	(10.926)	(6.196)	(5.708)
Costly Specialization	18.406^{**}	9.708^{*}	8.698^{*}
	(9.048)	(5.148)	(4.741)
Mixed Field	1.865	-1.104	2.969
	(9.048)	(5.148)	(4.741)
Order	-9.552	-4.333	-5.219
	(9.048)	(5.148)	(4.741)
Round 2	6.958	9.292^{*}	-2.333
	(10.002)	(5.632)	(5.191)
Round 3	10.813	14.000^{**}	-3.188
	(10.002)	(5.632)	(5.191)
Round 4	13.542	13.583^{**}	-0.042
	(10.002)	(5.632)	(5.191)
Num. obs.	192	192	192
Num. groups:	48	48	48

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

introduced by the spatiotemporal features of the environment did not prevent groups from achieving results close to the analytical benchmark estimate of what could be achieved under optimal harvesting.

Table 3.1 reports the results from the mixed effects model for R, which is further broken down into R_C for RGBY berry yields and R_S for silver berry yields. Note that for any observation $R = R_C + R_S$.

From Table 3.1 I reject the null hypothesis H_0^1 that costly specialization does not affect berry yields as measured by R. However, the null hypothesis H_0^2 cannot be rejected since the Mixed Field treatment does not affect berry yield. These estimates indicate that the costly specialization treatment leads to roughly 18 additional berries being harvested per round, which represents about 7.5% of the average yield across all sessions. This increase was not merely the result of better RGBY berry management, but is attributable to gains in both the number of RGBY and silver berries harvested.

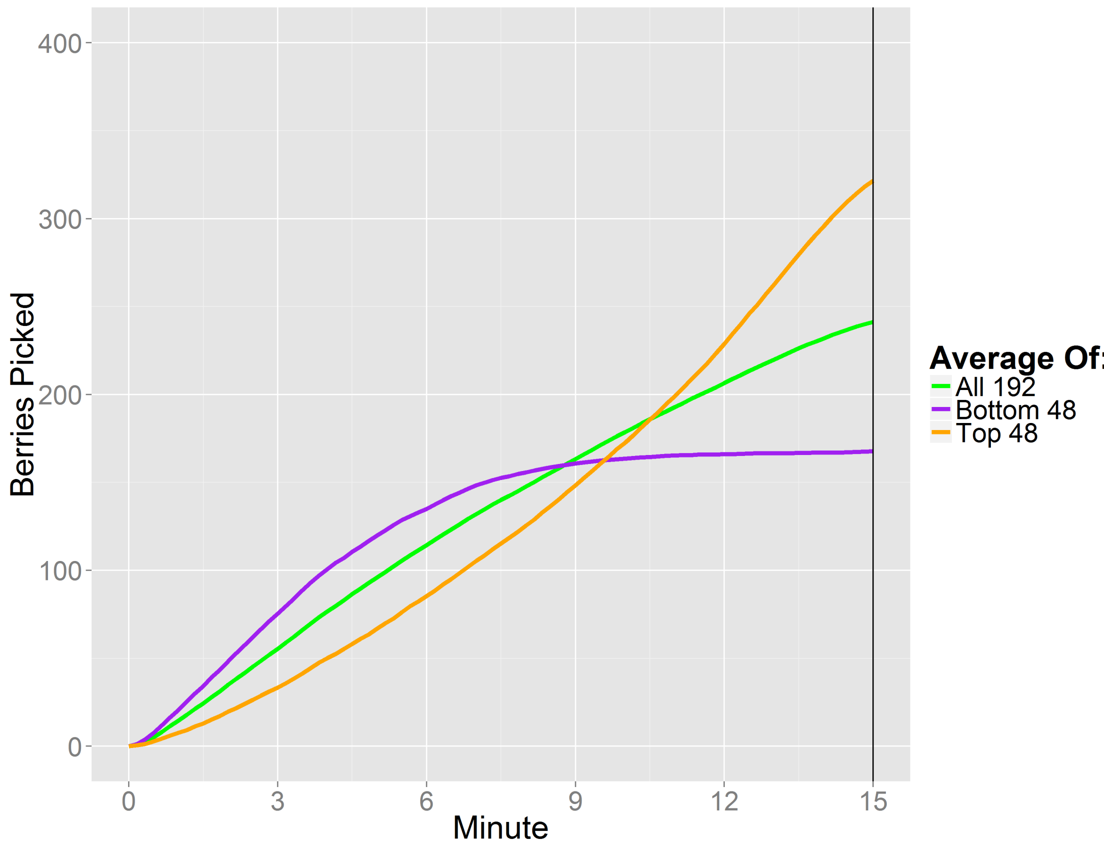

Figure 3.7: Total berry yield within rounds.

Across all observations, I find that many different strategies were used by different groups. Some groups tend to follow a strategy close to the Nash equilibrium prediction while other groups tend to follow a management strategy close to the planner strategy. These different strategies have clear theoretical implications for total berries harvested which are reflected in the data. Figure 3.7 shows the frequency at which berries were harvested in different groups. The "All 192" line represents the average of all 192 observations, while the "Top 48" and "Bottom 48" lines respectively represent the average of the top 48 and bottom 48 observations ranked by their total resource productivity R. Of particular note in this figure is the increasing convexity of R's evolution among groups that picked more berries overall, reflecting how restraint in harvesting is a necessary part of effective commons management in this environment.

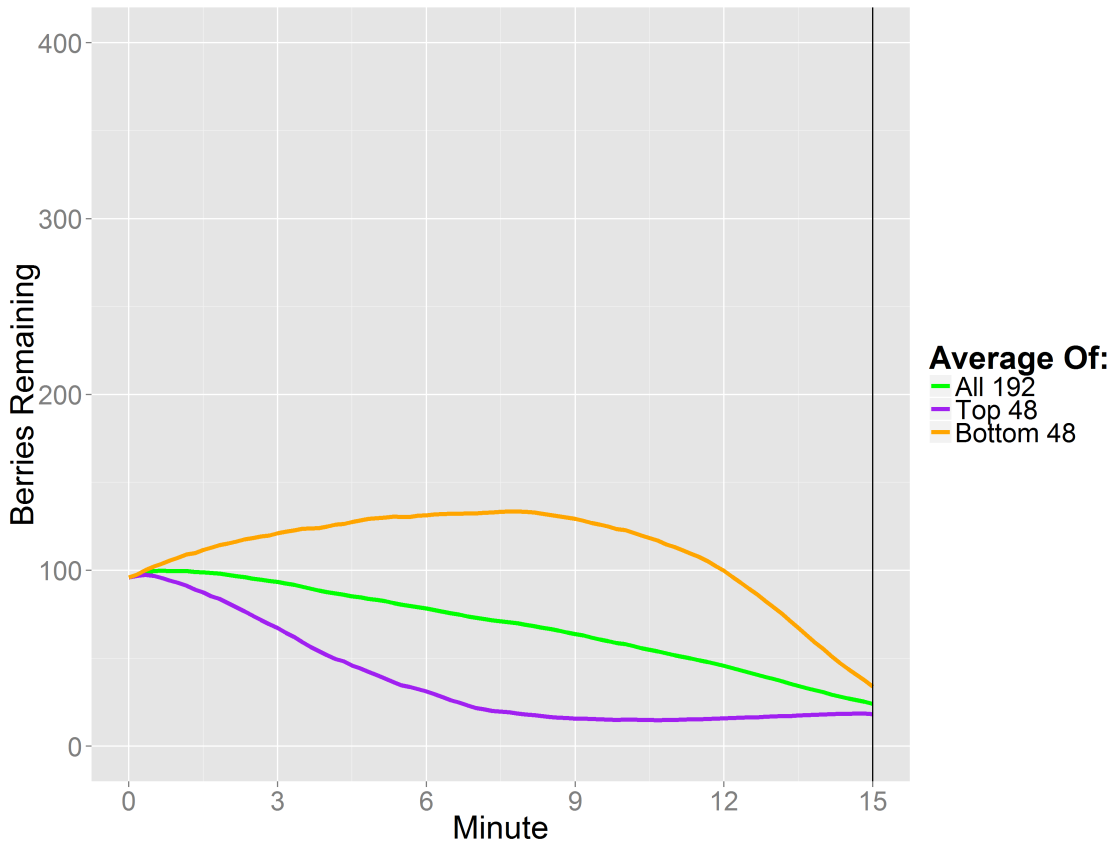

Figure 3.8: Total berries available within rounds.

Figure 3.8 shows how many of the berries were available on average within these subgroups throughout the round. Unsurprisingly, the bottom 48 observations exhibited a tendency to immediately reduce available berries through overharvesting. This overharvesting didn't drive the number of available berries to zero largely because subjects failed to check on some of the berry bushes as part of their clearing strategy. Of the 48 berry bushes, the median amount that had been completely cleared halfway through the round was 18. Figure 3.9 depicts what proportion of groups hit various clearing benchmarks throughout each 15-minute round. In addition to showing how many groups had cleared 18 at each point in time, this figure also shows how many had cleared a very low amount (8) and a very high amount (40). For example, 9 minutes into the period roughly 75% of groups had depleted

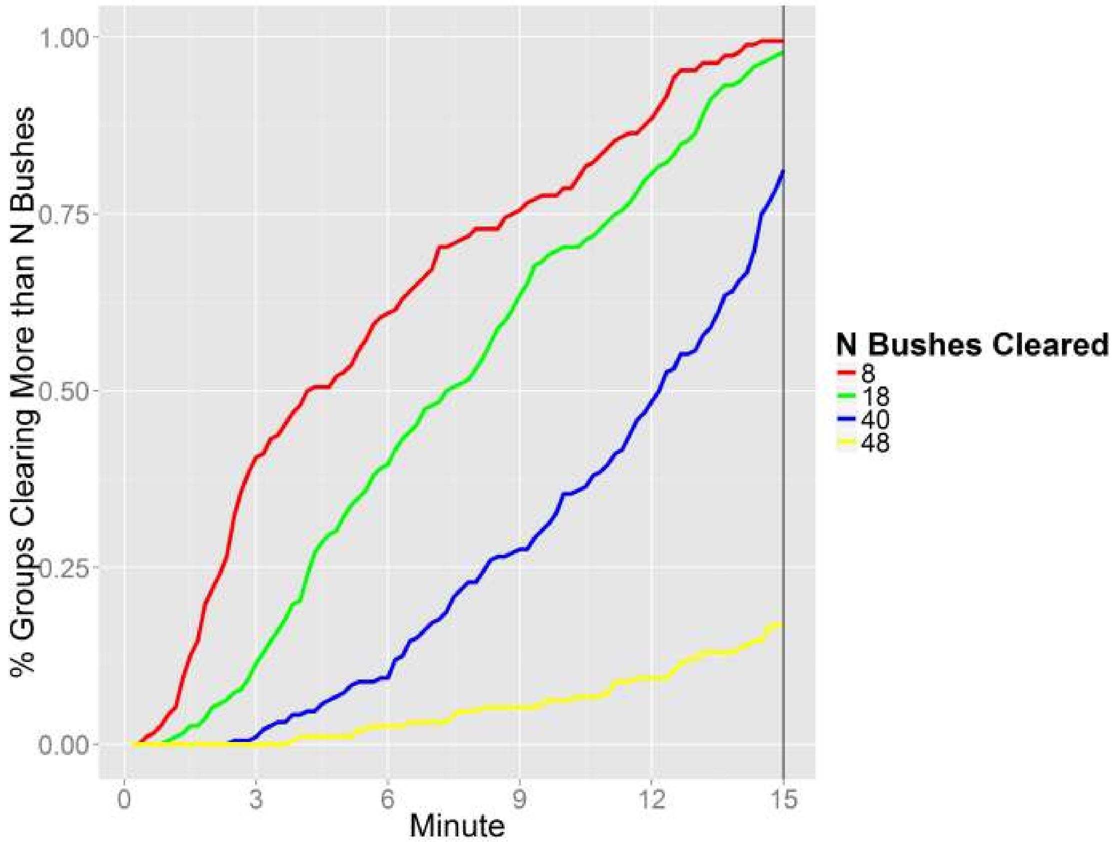

Figure 3.9: Bush depletion within rounds.

8 or more berry bushes, 64% had depleted 18 or more berry bushes, 28% had depleted 40 or more berry bushes, and 5% had depleted all 48 berry bushes. These figures illustrate the fact that even though many groups did not effectively manage the commons, these breakdowns in management oftentimes did not occur immediately but instead unfolded gradually over the duration of a given round.

3.3.2 Territorial Property Rights

H_1^2 hypothesizes that costly specialization will lead to stronger territorial property rights and therefore less territorial competition. In order to define the territoriality of subject groups, I used the minimum convex polygon approach that is common in the ecology literature (Aebischer, Robertson and Kenward, 1993; Mohr, 1947). For each subject i in each round of the experiment, a convex hull is defined for i containing all of the geographical coordinates corresponding to the normalized spatial locations of the bushes harvested by i. The environment's spatial coordinates are normalized so that each point in the environment can be represented as being in a $[0, 1] \times [0, 1]$ grid and refer to the convex hull assigned to subject i as i's "harvesting range" H_i. A harvesting density for any point (x, y) in the space defined by $D(x, y)$ can then be defined as:

$$D(x,y) = \sum_i d_i(x,y), \text{where } d(x,y) = \begin{cases} 1 & ,(x,y) \in H_i \\ 0 & ,\text{otherwise} \end{cases} \tag{3.6}$$

$D(x, y)$ counts the number of subjects whose territory hulls encompass a given point (x, y) in the environment. From this, the share of the environment's area that belongs to the harvesting range of exactly n subjects as M_n is defined as:

$$M_n = \int_0^1 \int_0^1 D_n(x,y)dxdy, \text{where } D_n(x,y) = \begin{cases} 1 & ,D(x,y) = n \\ 0 & ,\text{otherwise} \end{cases} \tag{3.7}$$

M_n is then used to define the final measure M:

$$M = \frac{M_1}{\sum_{n=1}^{4} M_n} \tag{3.8}$$

M represents the share of the environment that is uncontested, that is, in exactly one of the subject's harvesting range. It may be the case that subjects will establish different property rights regimes over silver berry bushes than over non-silver ones. Therefore I further define a set of subcategories for M that are generated by only defining convex hulls of subjects using certain subsets of the berries they picked:

- M will be the territory measure defined by taking the harvesting of all berry bushes into account.
- M_C will be the territory measure defined by taking only the harvesting of RGBY berry bushes into account.
- M_S will be the territory measure defined by taking only the harvesting of silver berry bushes into account.

The additional measures M_C and M_S will play a key role in examining whether specialization impacts property rights over the berry bushes of particular colors or if it impacts property rights over territorial regions as a whole.

Figure 3.10 provides a graph of the harvesting ranges generated for one group over its four round observations where (a), (b), (c) and (d) reflects rounds 1-4 respectively. This figure provides an illustrative example of the emergence and slight unraveling of territorial harvesting patterns within one group in the experiment.

Table 3.2 reports the results from the mixed effects model for the territoriality measures M, M_C, and M_S. This table provides strong support for the hypothesis that costly specialization increases the harvesting range of uncontested territories. Across all groups, the average territoriality score M was .42, meaning that 42% of the field that belonged to anyone's territory was uncontested. I estimate that the costly specialization treatment condition leads to an additional 13% of the environment belonging exclusively to one subject's territory, while the mixed field treatment condition had no effect on territoriality. Figure 3.11 shows a histogram of the main territoriality measure M_A across the specialization treatment conditions in order to visualize the full distribution of this measure. Of particular note is that of the 41 observations where $M_A > .6$, 30 of them were generated under the costly specialization treatment condition. Figure 3.12 helps visualize this effect by depicting a heat map detailing the average number of territory hulls each point in the

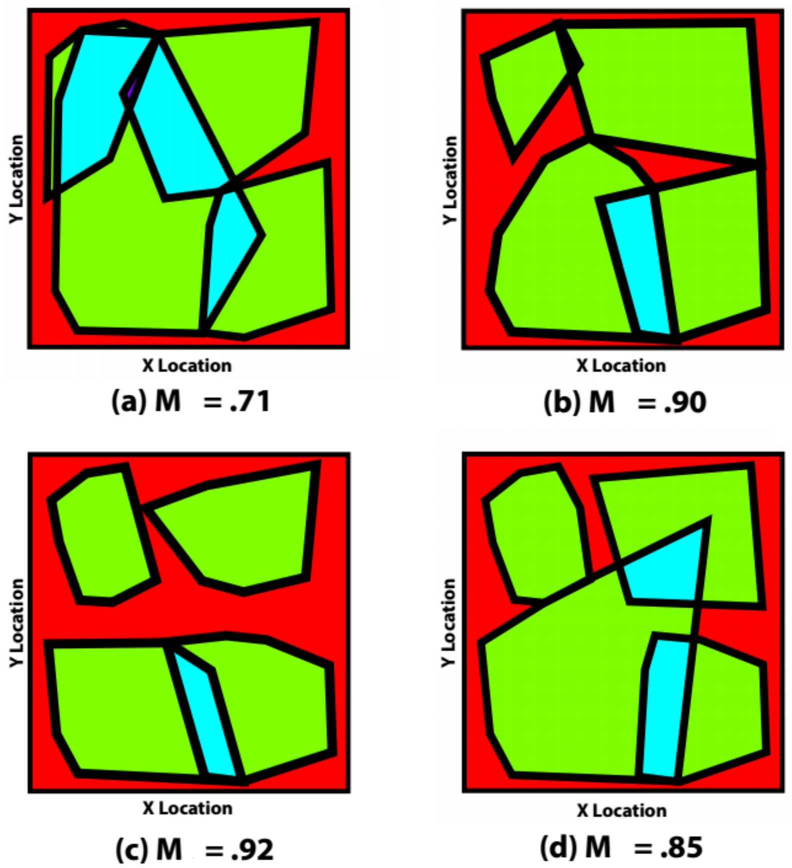

Figure 3.10: Spatial map depicting the locations and overlap of each subject's territory within a single group from rounds 1 to 4. Regions belonging to 0/1/2/3 subjects were colored as red/green/blue/purple respectively. Theuncontested territory measure (M) is the share of territory that belonged to at least one subject that belonged to exactly one subject.

environment belonged to in each specialization treatment condition - it is clear by inspection that there was less territorial overlap in the costly specialization treatment condition than the no cost specialization treatment condition. This effect persists across the different territoriality measures are employed. Importantly, the improvement of M_S in the costly specialization indicates that this treatment did not merely affect the management of the RGBY berry bushes, but instead affected the general management of uncontested spatial territories that included both RGBY and silver berry bushes. These results lead to the rejection of H_1^2 and H_0^2 and the acceptance of H_2^2.

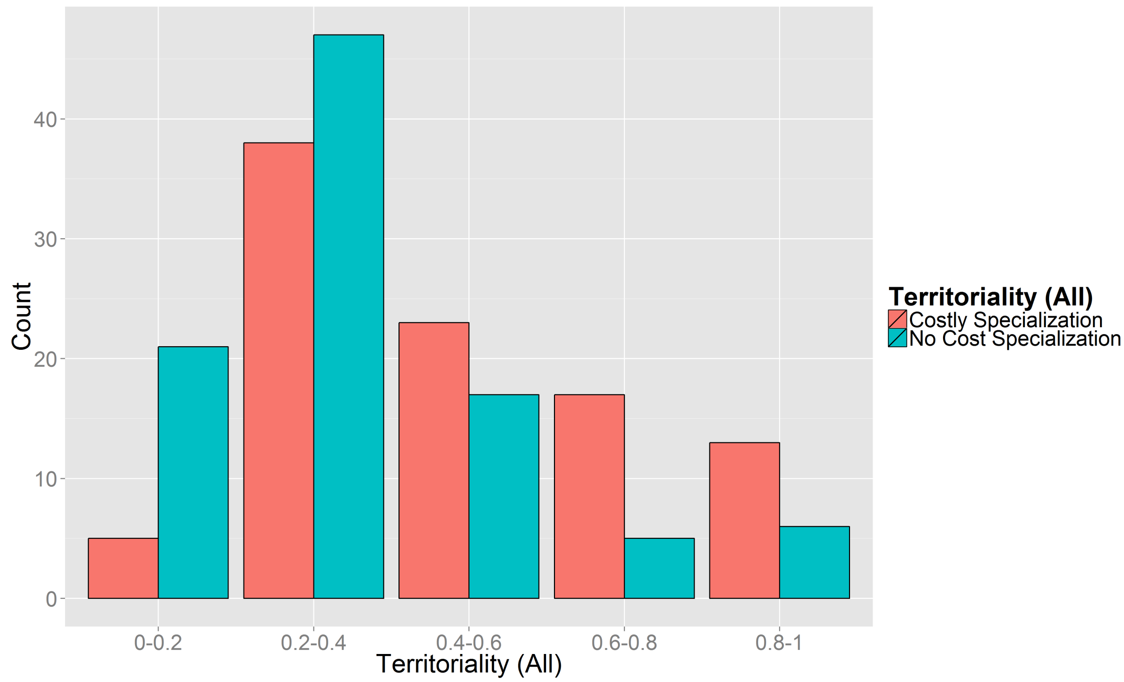

Figure 3.11: Frequency histogram of the exclusive territory share M_A in each group, by specialization condition.

Figure 3.13 plots the relationship between territoriality and resource productivity in the 192 rounds of data. The data can be divided rather cleanly into two clusters using a 2-means clustering procedure - one in the lower-left (Low (All)) within which there is a negative relationship between territory and productivity, and one in the upper-right (High (All)) that contains most of the highest-performing groups in the dataset. Of the 32 round observations in the High cluster, 24 of them are obtained in the costly specialization treatment condition. This disproportionate representation helps explain explains the positive correlation between M_A and R that exists within only the costly specialization treatment condition.

One reason to expect this clustering to occur is due to the fact that groups with poor territoriality may in fact still appear territorial given the definition of M. In particular, a group that simply clears the field as quickly as possible because of a lack of effective property rights, where each subject simply scatters in different directions and clears a portion of the field as quickly as possible, would be assigned a high M because each individual's harvesting

Table 3.2: Treatment Effects on Territoriality

	M	M_C	M_S
Intercept	0.221^{***}	0.398^{***}	0.357^{***}
	(0.051)	(0.058)	(0.049)
Costly Specialization	0.134^{***}	0.127^{**}	0.134^{***}
	(0.047)	(0.053)	(0.045)
Mixed Field	-0.009	0.019	-0.037
	(0.047)	(0.053)	(0.045)
Order	0.122^{***}	0.119^{**}	0.128^{***}
	(0.047)	(0.053)	(0.045)
Round 2	0.080^{***}	0.110^{***}	0.073^{**}
	(0.030)	(0.039)	(0.035)
Round 3	0.123^{***}	0.132^{***}	0.138^{***}
	(0.030)	(0.039)	(0.035)
Round 4	0.092^{***}	0.138^{***}	0.080^{**}
	(0.030)	(0.039)	(0.035)
Num. obs.	192	192	192
Num. groups:	48	48	48

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

range will be independent. This illustrates one manner in which M may overestimate the actual extent of territoriality in a group. For example, in the sole group for whome $M > .75$ and $R < 150$, more than 40 berry bushes were depleted by the time 5 minutes had passed.

To get an idea of how territoriality M correlates with resource productivity R in the groups who quickly deplete environmental resources, observations are once more dividied based on whether they have depleted more or less than the median number of bushes halfway through the round. I call groups that have depleted at least 18 bushes fast depleters and groups that have depleted less than this amount slow depleters. Splitting Figure 3.13 based on these groups results in Figures 3.14 and 3.15. Of the 32 observations in the previous "High (All)" cluster, only 5 of them came from fast-depleting observations.

The slow depletion groups are those who are most likely to have successfully implemented some sort of management rule that prevents the overharvesting of the commons. Looking at Figure 3.14, we can see what appears to be a clustering of observations into two different regions. 2-means clustering is again used to designate observations as belonging

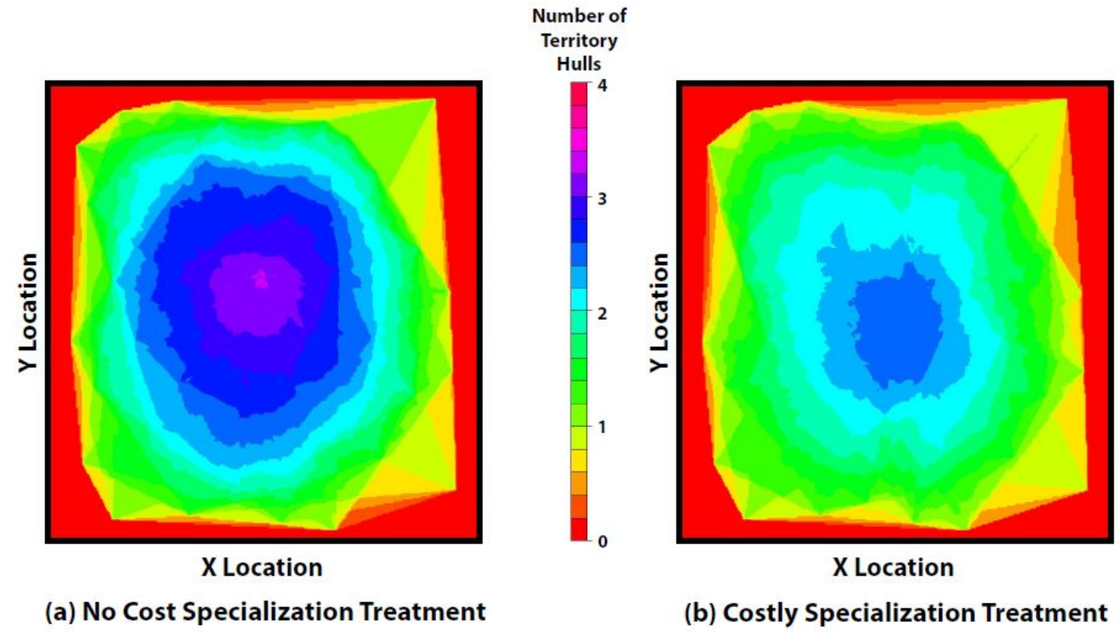

Figure 3.12: Heatmap of the average number of subjects who had each point in the environment within their territory.

to either a "Low (Slow)" cluster or a "High (Slow)" cluster. Groups in the "Low (Slow)" cluster - that is, groups that exhibited above-median commons management but not high territoriality - can be considered to have established more-communal systems of property rights management than those in the "High (Slow)" cluster. As a robustness test of the hypothesis that costly specialization induces groups to adopt more-territorial management rules, I use a fixed-effects probit regression to investigate whether the costly specialization treatment effect predicts whether an individual observation in the slow depletion subset of observations will belong to the "High (Slow)" cluster. The results of this estimation can be observed in Table 3.3, which shows that costly specialization does indeed predict that low-depletion groups will establish less-communal management regimes.

3.3.3 Specialization and Coordination

Under the Costly Specialization treatment condition, the maximization of income by group members involves each subject choosing a different color to specialize in. Despite subjects

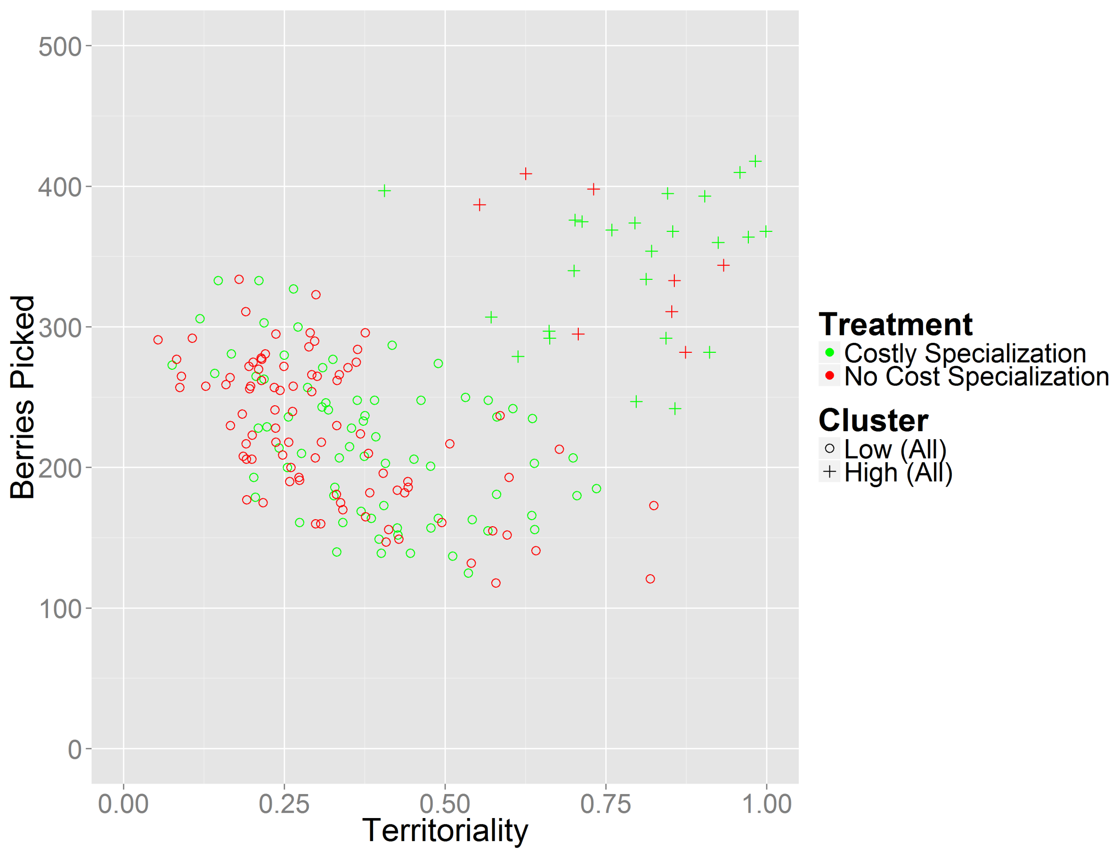

Figure 3.13: Territoriality (M) vs. Resource Productivity (R), with clusters assigned via k-means clustering.

being able to monitor one another and communicate in order to avoid choosing the same colors, perfect coordination (i.e. 4 unique specializations) is observed in only 28 out of the 96 observations under this treatment condition. However, as shown in Table 3.4, these 28 groups performed better than the others both in terms of overall resource productivity and in terms of territoriality. Territoriality does not vary monotonically with the number of unique specializations within observations, but hovers at an average of roughly .4 except in the case where all subjects have chosen unique specializations, in which case it averages .7. This increase may indicate that any amount of miscoordination on specialization could inhibit the formation of private property rights, although the direction of causality for this effect cannot be clearly identified.

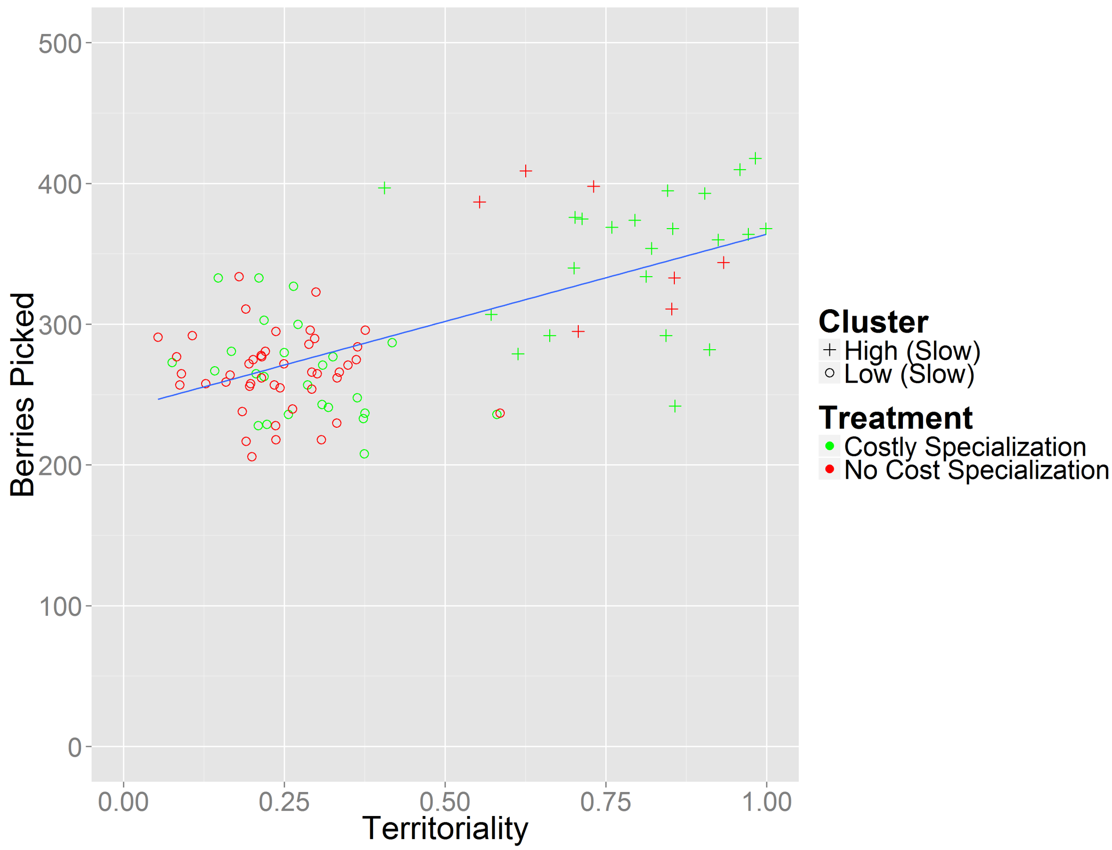

Figure 3.14: Territoriality (M) vs. Resource Productivity (R): Slow Depletion Groups, with clusters assigned via k-means clustering.

3.3.4 Communication, Punishment and Trade

Table 3.5 shows estimates of how the treatment conditions and other covariates affected the aggregate volume of different behaviors in this experiment. The first column of this table pertains to the total communications volume measured in chat messages sent between subjects. Subjects averaged 23 messages per group per round, and 84% of these messages were sent at the focal meeting campfire in the middle of each field. Similar levels of communication were seen in both successful and unsuccessful groups - observations where punishment was present had a mean harvest of 221.94 berries with a standard deviation of 57.80, while observations without punishment had a mean harvest of 246.85 berries with

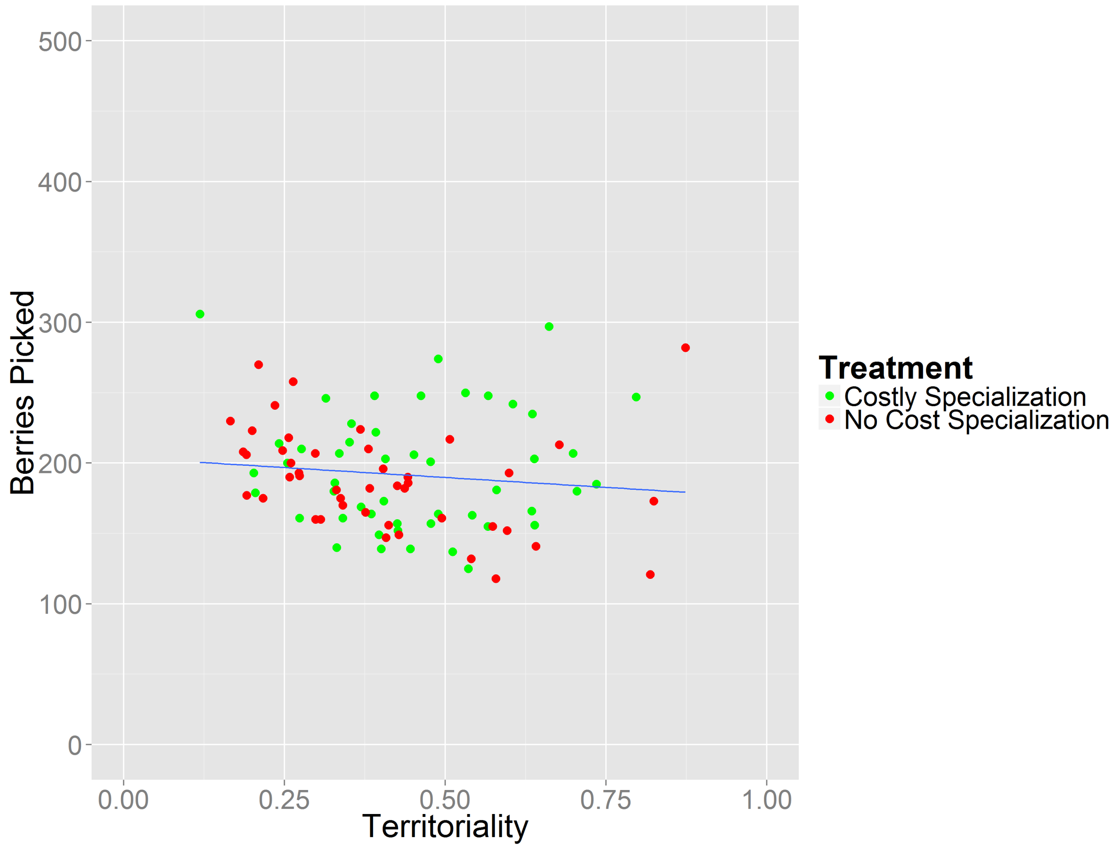

Figure 3.15: Territoriality (M) vs. Resource Productivity (R): Fast Depletion Groups

a standard deviation of 69.42. Figure 3.16 plots the cumulative distribution of chat messages per observation by experimental round. Communication served a variety of purposes in this experiment, including establishing and maintaining territories and property rights, conveying threats, and building trust. The chat data included cases where groups especially talked about setting the various parameters of management rules $g(c, \tau)$. A full set of chat logs from each group in this dataset can be found in Appendix C of this paper's online supplementary materials.[8] Here are a few illustrative examples:

Example 1: Assertion of property rights and discussion of punishment within a group, and a proposal to set the harvesting cutoff (c) to 4

[8]These materials can be found at `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2513369`.

Table 3.3: Treatment effects on private rights management among low-depletion groups

	Slow Depletion / High Territoriality
Intercept	-4.481^{***}
	(1.627)
Costly Specialization	2.807^{**}
	(1.315)
Mixed Field	0.766
	(1.264)
Order	2.769^{*}
	(1.592)
Round 2	1.077
	(0.936)
Round 3	1.046
	(0.945)
Round 4	0.284
	(1.062)
Num. obs.	96
Num. groups:	24

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3.4: Summary statistics on coordination, with mean berry harvest (R) and territoriality (M) measures

Unique Specializations	Total Count	Round 1 / 2 / 3 / 4 Count	$\bar{R}$	$\bar{M}$
1	4	2 / 0 / 0 / 2	172.50	.4140
2	22	4 / 8 / 4 / 6	206.23	.3764
3	42	14 / 11 / 10 / 7	238.26	.4047
4	28	4 / 5 / 10 / 9	309.96	.7021
Total	96	24 / 24 / 24 / 24		

Oregon Wirsing: stop

Oregon Wirsing: this is my territory

Oregon Wirsing: stop

Mississippi Halsey: stop punishing me...

Oregon Wirsing: stop picking on my territory

Oregon Wirsing: i will use take blue first

Table 3.5: Treatment Effects on behavior

	Messages Sent	Punishment Tokens Used	Seeds Transferred
Intercept	19.792^{***}	15.927	9.974^{***}
	(5.310)	(12.382)	(3.525)
Costly Specialization	1.437	-9.604	-2.760
	(4.682)	(9.360)	(2.955)
Mixed Field	1.146	-6.729	-2.573
	(4.682)	(9.360)	(2.955)
Order	10.083^{**}	2.771	-2.865
	(4.682)	(9.360)	(2.955)
Round 2	-5.667	3.646	-2.562
	(4.092)	(13.237)	(3.141)
Round 3	-3.333	11.625	1.063
	(4.092)	(13.237)	(3.141)
Round 4	-2.458	21.229	1.063
	(4.092)	(13.237)	(3.141)
Num. obs.	192	192	192
Num. groups:	48	48	48

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Mississippi Halsey: stop

Mississippi Halsey: it's self destructive

Oregon Wirsing: wait

Oregon Wirsing: lets talk

Oregon Wirsing: just for a sec

Mississippi Halsey: stop being a jerk

Oregon Wirsing: just take the blue ones

Oregon Wirsing: don't pick on green

Oregon Wirsing: i will be ata a loss

Oregon Wirsing: and keep the quantity to 4

Oregon Wirsing: lot of regeneration at 4

Mississippi Halsey: OK. No more punishment tokens. You've hurt both our earnings because you decided to be a mean spirited jerk

Oregon Wirsing: ok

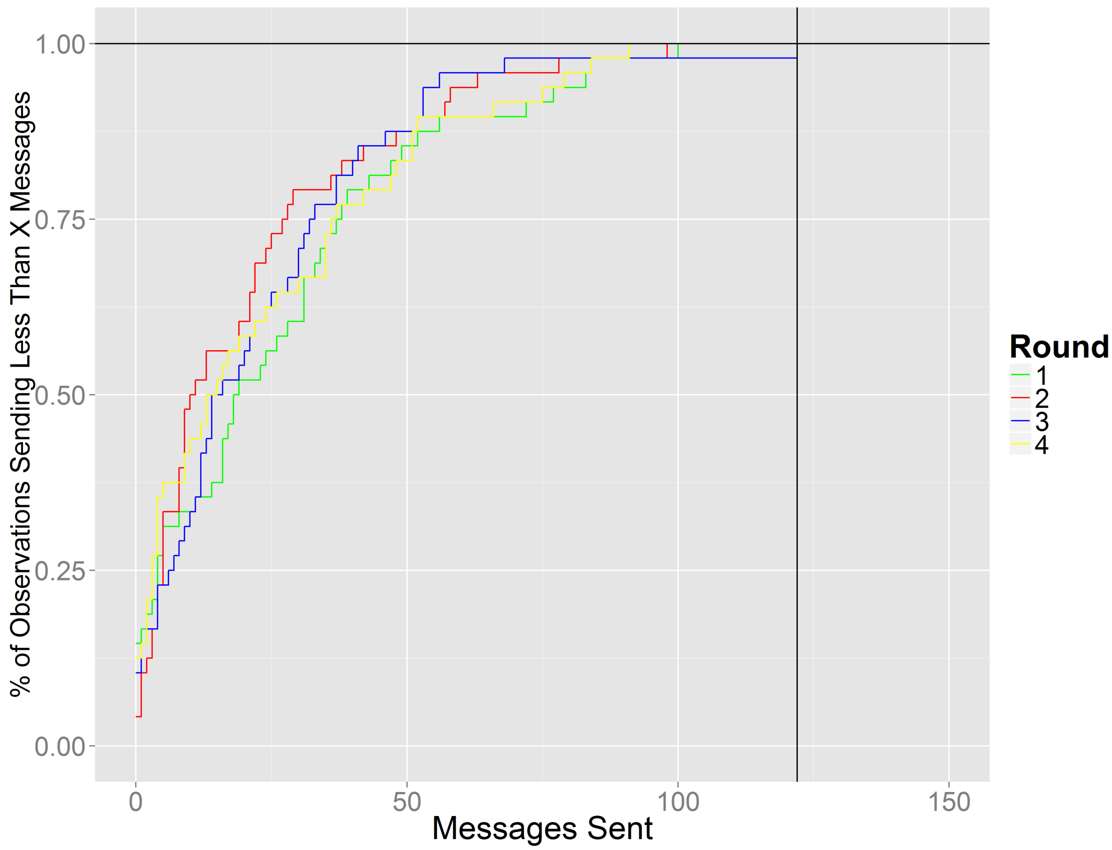

Figure 3.16: Cumulative distribution of chat volume, by round

Oregon Wirsing: everyone lets just plan

Connecticut Nowles: i'll be yellow next round

Oregon Wirsing: i will take green

Connecticut Nowles: u want blue missi

Mississippi Halsey: If you do one thing Oregon doesn't like, he'll take away all your earnings

Oregon Wirsing: i 'm at a loss here

Oregon Wirsing: becoz of u

Mississippi Halsey: you earned it

Example 2: A discussion of how to manage the clearing time of silver

berries

Hampshire Sleydon: and kill all silvers at 2 mins left

Maine Falmer: works for me

Hampshire Sleydon: start harvesting all

Dakota Mosely: okk

Hampshire Sleydon: even from other territory

Dakota Mosely: okk

Hampshire Sleydon: till then same rules

Dakota Mosely: nobody attacks the others trees

Maine Falmer: so until 2, only your colors, and then free for all kill

Hampshire Sleydon: all kill only

Hampshire Sleydon: applies to silvers

Maine Falmer: got it

Hampshire Sleydon: basically you'll spend time

Hampshire Sleydon: killing you own color

Hampshire Sleydon: as well

Maine Falmer: true

Maine Falmer: same colors and plan?

Dakota Mosely: seemed to wrk fine

Hampshire Sleydon: yupp

Hampshire Sleydon: much better

Maine Falmer: cool

Dakota Mosely: yeah

Hampshire Sleydon: also get as many silvers as you can towards the end

Example 3: Public punishment

Wisconsin Murfin: you're a jerk kansas

Wisconsin Murfin: dont deplete

Wyoming Harvy: nice jump

Wisconsin Murfin: i hate kansas

Wyoming Harvy: what they do this time?

Wisconsin Murfin: nothing new i'm just assuming they depleted all the berries

Wyoming Harvy: i punished kansas with 200 tokens, sooo im sure they lost all their money

Wisconsin Murfin: haha yeah i'm glad you did

These chat excerpts indicate that at least some subjects were willing to use the punishment tokens made available to them. Overall, punishment was used only 54 times in 33 of the 192 observations. When a subject chose to punish, the average amount of tokens used was roughly 65. Although most punishment incidents occured under the no cost specialization treatment condition, this effect was statistically insignificant. Additionally, the use of punishment was observed in groups that were both successful and unsuccessful in managing the commons. Note that this measure only includes actual observed punishment and not factors such as threats of punishment or subjects acting out of a fear of punishment. Even in groups where no punishment was observed, the capacity for punishment to happen may have molded subject behavior.

The use of transfers was as rare as the use of punishment, being seen in 35 out of 192 observations in total. Most of these transfers were reciprocated by transfers back from the other subject, indicating that they constituted a form of social grooming whereby exchanges were being made that did not actually leave either subject materially better off. Among these 35 observations, only 13 involved cases where a net transfer of resources from any subject to another subject was made.

3.4 Concusion

This experiment demonstrates that by credibly and publicly committing to only receiving an enhanced yield from one variety of a spatially-clustered common resource, groups are able to more effectively avoid the tragedy of the commons. The improved management observed under the costly specialization treatment condition coincides with the establishment of stronger territorial harvesting patterns, indicating that this treatment is affecting the character of the norms governing the management of the CPR by pushing groups towards stronger systems of territorial property rights.

This result builds on our existing knowledge of how groups manage common resources. By observing these results in an ecologically-robust environment where subjects already have access to many of the tools that have been previously-implicated in good commons management, it is shown that further incremental progress can be made in achieving a deeper understanding of the determinants of effective common management regimes through the implementation of increasingly robust environments that allow for subjects to solve these dilemmas in naturalistic ways. As seen in real-world communities such as Wikipedia, costly and publicly-observable specializations can provide a natural mechanism for generating informal private property rights. This result raises several further questions: What sorts of communities can this sort of specialization exist in? How can communities enable members to signal specializations in order to define territorial boundaries? Under what conditions might it be desirable to induce specializations via policy in order to promote the establishment of private property rights? I leave these questions to further research.

Chapter 4: Trust Networks

In this study I investigate the emergence of specialization in production and long-distance exchange within a limited property rights environment in an experimental setting. In the field, when property rights are either weak or non-existent exchange relationships tend to be characterized by their personal nature. To borrow the typology outlined in North (1990, 1991), this sort personal exchange manifests through repeat interactions between agents that have formed relationships based on trust, reciprocity, and established reputations. However, the development of these mechanisms may also be inhibited through rent-seeking behaviors (DeScioli and Wilson, 2011; Leeson, 2007; Olson, 1993). Although these sorts of personal exchange relationships have been important and ubiquitous both historically and in the present, the precise conditions under which effective instititutions of personal and impersonal exchange arise has not been extensively investigated within a controlled experimental context. This paper aims to contribute to our knowledge in this area by investigating the particular question of how the characteristics of networks of trust and trade within and between groups respond to the joint presence of both an excess supply of labor and weak property rights over productive resources.

The main finding of this experiment is that the presence of excess labor inhibits the effective formation of networks and trust and trade within groups. This result is observed within a spatial microeconomic system where agents are able to engage in specialized production and positive-sum exchange; however, access to production opportunities is governed by a system of squatters' rights whereby once an agent leaves an area used for production another may claim exclusive control over it. It has been shown in the field that these sorts of insecure property rights can lead agents to substitute their labor expenditures inefficiently towards home production (Clay, 2006; Field, 2007). In this environment, "excess labor" connotates the fact that the number of agents within a group may exceed the number of productive areas that can simultaneously use them. The novel result that is demonstrated here is that the negative impact of these weak property rights is modulated by the extent to which this excess labor exists. On net, rather than encouraging the establishment of a specialization of roles whereby some agents engage primarily in production and others engage in exchange, the presence of an excess labor supply leads to a greater degree of autarky within groups which negatively impacts the overall extent of exchange that occurs. This result is corroborated through an examination of the trade networks that arise (or do not arise) under the different treatment conditions that vary this excess supply of labor.

This research contributes to a broader pre-existing literature concerning the property rights systems that foster the development of effective trust and trade relationships between specialized parties. Greif (1993, 2006) extensively documents the trust-based community responsibility system that existed among medieval traders in order to facilitate the historical development of long-distance exchange. Additional research such as Fafchamps et al. (1998), Fafchamps and Minten (1999) and Karlan et al. (2009) provide contemporary examples of exchange networks that rely on personal exchange mechanisms in the absence of secure property rights. This experiment intersects this work with the aforementioned empirical literature concerning the impacts of insecure property rights over productive resources.

This experiment does not represent the first contribution of a laboratory experiment to questions concerning the emergence of trust and trade within groups. In Kimbrough, Smith and Wilson (2008, 2010) and Crockett, Smith and Wilson (2009), the emergence of personal and impersonal patterns of production and exchange is explored in an economy where subjects could choose to produce two or three different goods and discover mutually advantageous exchange opportunities. Within these economies, each subject was given a comparative advantage in the production of a single sort of good and the key social dilemma that was presented concerned whether subject groups could discover and converge to a wealth-maximizing pattern of production and exchange. These interactions took place using the experimental interface depicted in Figure 4.1, which allowed for subjects to communicate via a shared chat window as well as transfer resources from one another by moving around

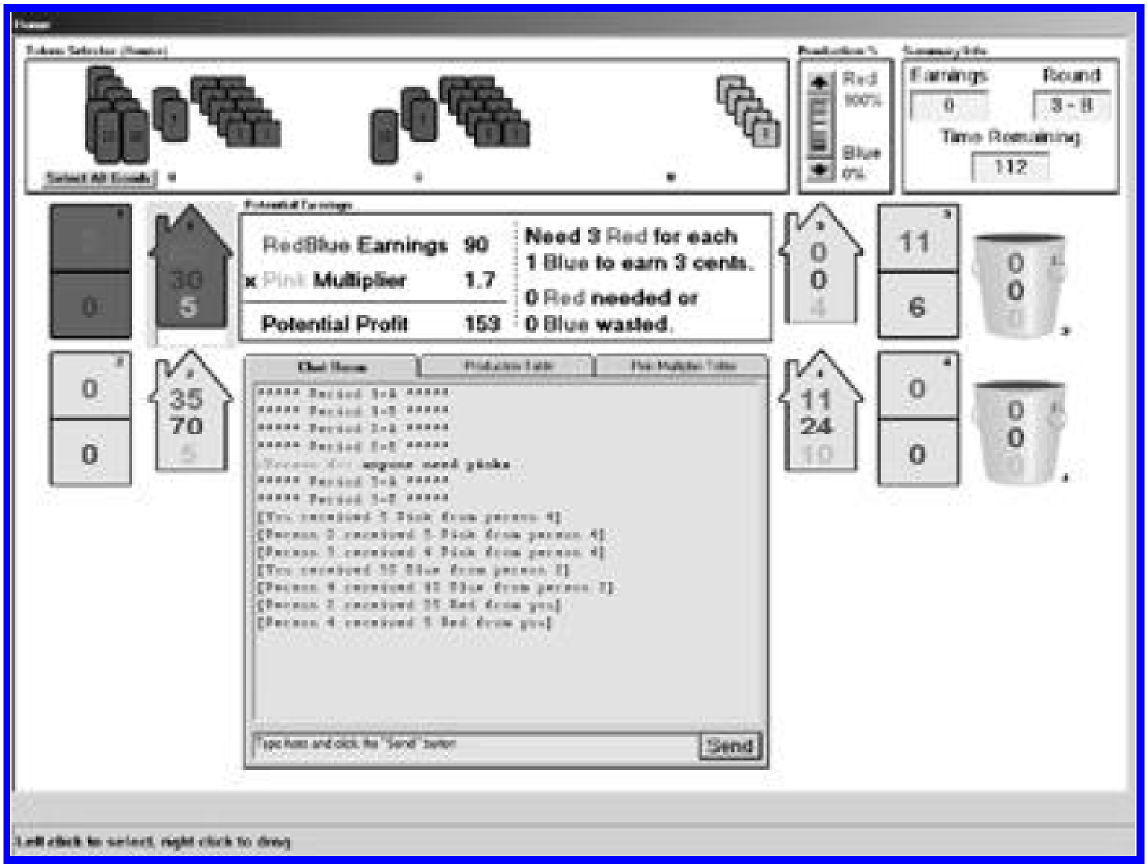

Figure 4.1: Subject view of the Kimbrough, Smith and Wilson (2008) interface.

representative icons within the graphical user interface. Kimbrough and Wilson (2011) subsequently expanded on this design by investigating the social characteristics of groups that developed successful exchange intermediaries.

The experimental design that will be detailed in the next section follows this broad template outlined in Kimbrough, Smith and Wilson (2008, 2010). In the experiment, subjects are divided into three different "villages". Each village contains a number of facilities known as "harvesters" that can be used to produce different commodities that can be combined together to create higher-order goods. This complementarity between different resources provides incentives for subjects to either use different harvesters (each harvester produces a single type of good) or to engage in exchange. Furthermore, complementarities between

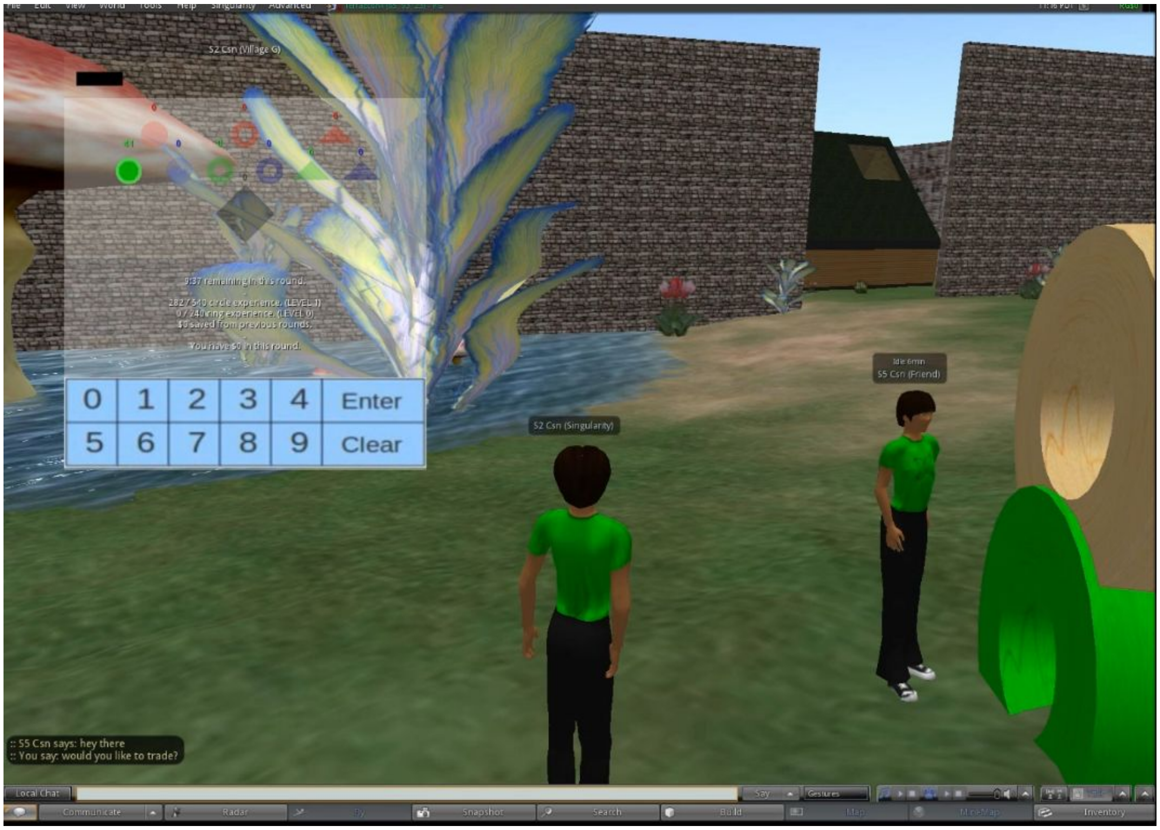

Figure 4.2: Two subjects meeting at a green ring harvester within the experimental interface.

goods from different villages exist, further incentivizing subjects to develop networks of long-distance exchange where goods are traded both within and between different village groups.

This experiment took place in a virtual spatial environment where subjects controlled avatars that interacted both with different experimental objects such as the harvesters as well as one another. This environment was instantiated within the platform of OpenSim, which provided several built-in features for constructing avatar-based experimental environments. Importantly, however, the virtual world utilized in this design allows for a naturalistic implementation of the motivating spatiotemporal dilemma concerning the emergence of trade in the presence of excess labor. This is due to the fact that avatars must navigate the spatial environment in order to find exchange partners, and in so doing may face a tradeoff in having to choose between being bound to one location in order to engage in production and being able to travel through the environment in order to find trading partners. This virtual environment is depicted in Figure 4.2, which shows two subjects meeting at a harvester.

Despite the spatiotemporal dynamics of this virtual environment providing subjects with a fairly robust strategy set, the core social dilemma faced by subjects here should be somewhat familiar. In this experiment, the problem of producing higher-order goods through exchange is implemented in a fashion that is meant to replicate the payoffs of the workhorse trust game in Berg, Dickhaut and McCabe (1995). Trade in this experiment is essentially a two-stage dictator game where a subject i must make a unilateral transfer of goods in real time to subject j in the hopes of reciprocation. While this may be a relatively low-risk transaction in a direct exchange[1], this may prove to pose a more difficult problem when it comes to long-distance exchange using intermediaries. Under this scenario, if a third subject k proposes to relay goods from i to j and then back to i (keeping some share of the proceeds to herself), then this process will likely expose i to much more risk than direct exchange. Prior research has indicated that the transformation of a direct trust relationship into an intermediated "pass-through" one can have deleterious effects on trust behavior, particularly when the actions of the intermediary and trustee cannot be observed (Kimbrough and Wilson, 2011; Ostrom and Walker, 2003; Rietz et al., 2013).

4.1 Experiment and Hypotheses

4.1.1 Harvesting Task

This experiment is built around a repeated 15-minute production and exchange task. In this task, a group of twelve subjects is divided into three villages each identified by a unique color (red, green, or blue) and each comprised of four subjects. Additionally, each village possesses a number of harvesters that can be used by subjects belonging to that village in

[1] Subjects are free to salami slice any sort of exchange they are trying to execute in order to limit their exposure to counterparty defection.

order to acquire resources of two different basic types called circles and rings. Each harvester in this environment is either a circle harvester or a ring harvester and produces goods of the according type. These resources are also produced in a given village's color, so that the harvesters in the red village can be used to produce red rings and circles, the harvesters in the green village can be used to produce green rings and circles, and the harvesters in the blue village can be used to produce blue rings and circles.

Each harvester can only be used by one subject at a time, and once a subject is using a harvester they will continue to have the exclusive ability to gather resources from it until they leave it to use a different harvester or find trade partners. This feature of harvesters allows subjects to only exert squatters' rights over them that expire when the subject leaves a harvester's proximity. While a subject is using a harvester it will produce goods at a stochastic rate, with there being a fixed probability of a harvester providing a yield every 3 seconds. For circle-producing harvesters, this probability is fixed to .3 while for ring-producing harvesters it is .15. This means that in expectation a circle harvester will produce goods 90 times over a 15-minute period while a ring harvester will produce goods 45 times over a 15-minute period. Because of the Poisson nature of this production process, the variance in the rate of production for circle and ring harvesters is likewise 90 and 45, respectively, over a 15-minute interval.

By default, each time a harvester produces a yield the subject who is using that harvester will earn a single ring or circle. However, subjects can become specialized in the harvesting of particular types of goods by using the harvester for a sufficient amount of time during a round. If a subject spends 4 minutes of the round using a harvester of a given type, they will receive two units of that good at a time when a yield is produced. Furthermore, if a subject spends 9 minutes of the round using a harvester of a given type, they will receive four units of that good at a time when a yield is produced.[2] This specialization process is key to how exchange is incentivized in this environment: Notably, a subject i and a subject

[2]These time investments need not be made within a continuous interval, however. A subject can harvest for several minutes, leave their harvester, and then return in order to continue where they left off in progressing towards specialization.

j who respectively spend a full 15 minutes at a circle harvester and a ring harvester will produce an expected joint yield of 760 circles and 380 rings, while a subject i and a subject j who split their time evenly between both harvester types will produce an expected joint yield of 440 circles and 220 rings.

Subjects attain income in this task by acquiring the resources available in this environment and holding them at the end of the 15-minute period. For every subject, each circle or ring held at the end of the round is worth 10 experimental dollars regardless of its color. However, these resources can also be combined with one another in order to create more-valuable goods. A ring and a circle of the same color can be combined into a triangle of that color, which is worth 40 experimental dollars at the end of the round. Two triangles of different colors can be combined to create a diamond, which is worth 160 experimental dollars at the end of the round. Additionally, subjects could at any time liquidate resources they were holding into experimental dollars that could be transferred among one another as a sort of direct utility transfer.

Ultimately, a subject's payoff from the harvesting task can be represented as:

$$U = \Pi + 160 * D + \sum_{c \in \{r,g,b\}} 10 * (R_c + C_c) + 40 * T_c \tag{4.1}$$

Where:

- Π is the amount of liquidated experimental dollars held by the subject.
- D is the number of diamonds held by the subject.
- R_c is the number of rings of color c held by the subject.
- C_c is the number of circles of color c held by the subject.
- T_c is the number of triangles of color c held by the subject.
- $\{r, g, b\}$ are the set of possible colors associated with rings, circles, and triangles.

The main treatment condition that is used in this experiment manipulates the excess labor supply in each village by altering the number of harvesters available for use. Under the 2-harvester (2H) treatment condition, each village contains a single circle-producing harvester and a single ring-producing harvester. Under the 3-harvester (3H) treatment condition, each village contains one circle-producing harvester and two ring-producing harvesters. Under the 4-harvester (4H) treatment condition, each village contains two circle-producing harvesters and two ring-producing harvesters. A crucial implication of this treatment variable is that under the 2H and 3H treatment conditions, there are more subjects within a village than there are harvesters. It is this modulation of harvester quantity that allows for the investigation of this experiment's main questions concerning how networks of trust and trade are impacted by the presence of this excess supply of labor.

4.1.2 Mechanisms of Exchange

There are two important observations to be made about this process through which higher-order goods are made: Firstly, while triangles can be acquired through both production and exchange within a single village, diamonds can only be acquired through intervillage exchange. There is no way for a subject belonging to a village of one color to acquire triangles of a different color without having them transferred from a villager belonging to that village. Secondly, the ratio of values between goods of different orders were chosen to reflect the ratio of values used in the canonical trust game outlined in Berg, Dickhaut and McCabe (1995).

In this environment, exchange exists as a sequence of unilateral transfers between subjects. If subject i wishes to trade 10 blue rings for subject j's 10 blue circles, then one subject will simply have to transfer some amount of resources to the other and hope that his partner reciprocates. Although subjects in our experiment were allowed to discuss any terms of exchange with their potential partners, there was no way for subjects to formally bind themselves to fulfill these terms. Therefore subject i's exchange with j can be seen as an act of trust in both j's intention and ability to fulfill the terms of trade. If long-distance trade emerges that uses intermediaries, then this exchange relationship may involve additional parties with whom relationships of trust and trustworthiness would have to develop.

However, there is an additional social mechanism through which triangles can be acquired: Subjects can share the usage of different harvester types so that they produce both circles and rings individually, and then combine them into triangles. Due to the scarcity of harvesters, particularly in the 2H and 3H treatment conditions, this may require active coordination between the users of different harvesters to exchange control of the harvesters; however, it does not require explicit trust, as there is not an individual incentive to defect on an agreement to swap harvesters in the same way that there is an individual incentive to defect on a trade agreement, since there is no way for an individual to acquire a new harvester without giving up the use of their old one.

4.1.3 Mechanisms of Social Interaction

Within the experimental environment, interactions between subjects will generally involve the management of usage rights over harvesters, as well as finding and establishing relationships of trust and reciprocity with potential exchange partners. The spatiotemporal features of the experimental environment introduce several naturalistic frictions associated with accomplishing these tasks. Importantly, however, the presence of these frictions are crucial to examining hypotheses concerning why one should expect patterns of exchange to be negatively impacted when excess labor supply is present.

One obvious friction associated with the experimental environment concerns the fact that subjects can only interact when within close proximity of one another in the spatial environment. As shown in Figure 4.3, the virtual environment was spatially divided into three different villages. Each village contains between two and four harvesters that could be used by subjects, with ring harvesters at one edge of the village and circle harvesters at the other. Because of the spatiotemporal features of this environment, the ability of subjects to monitor their surroundings and interact with one another was localized in order to reflect naturalistic constraints on these behaviors. Consequently, subjects had a limited

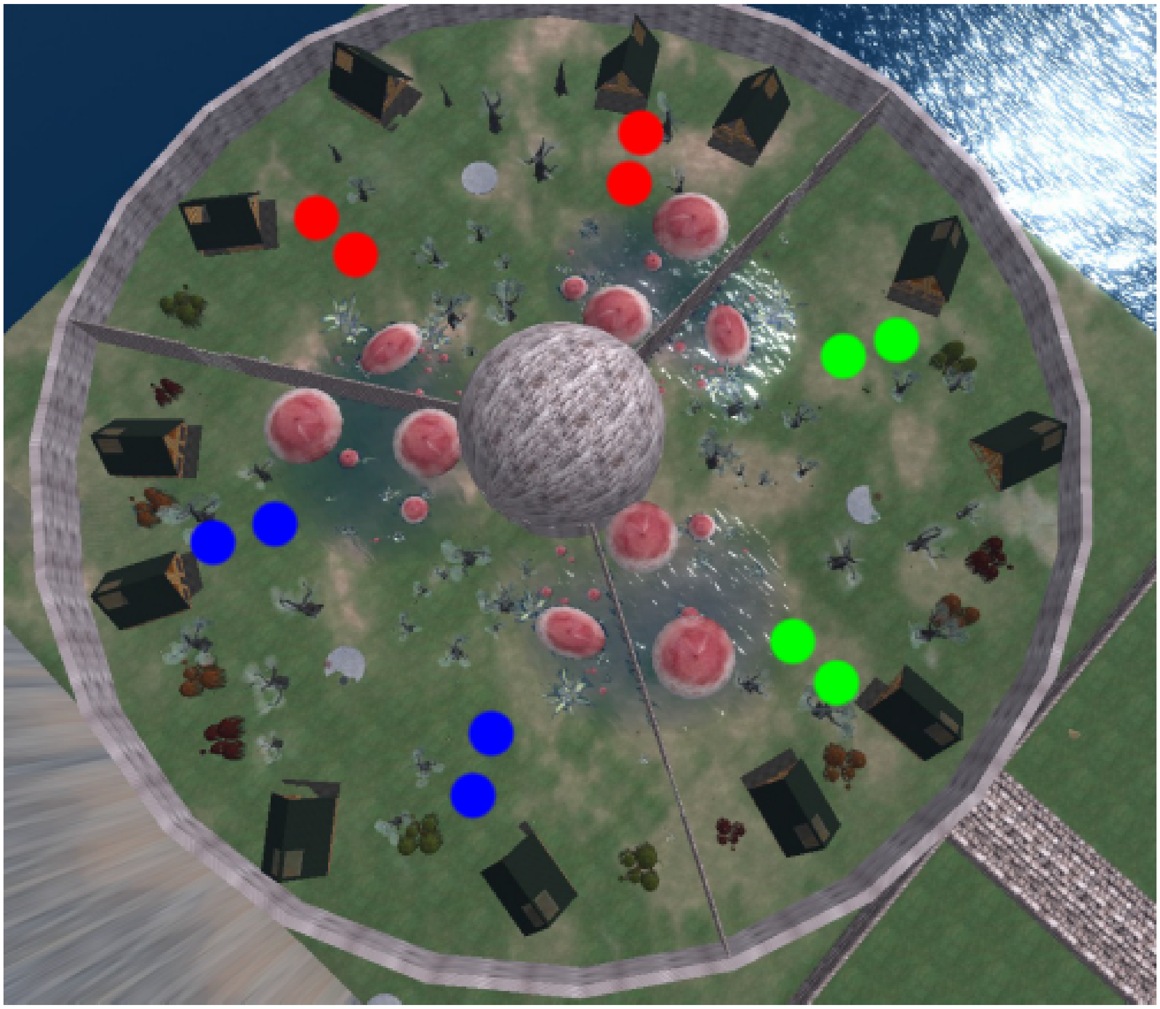

Figure 4.3: Overhead view of experimental environment, with color-coded highlighting of harvester locations. Traversal of a single village by an avatar takes roughly 30 seconds.

ability to see the status of the other harvesters and the positions of other avatars in the experiment. Under the 4H treatment condition a subject standing at a ring harvester would be able to see the other ring harvester and communicate with anyone who is using it, but would be unable to see the status of the circle harvesters, let alone what is going on in other villages. Traveling from the ring harvesters to the circle harvesters would take an avatar roughly 20 seconds, which meant that if a subject i was producing circles and wished to engage in exchange with a subject j producing rings, at least one of them would have to leave the harvester in order to find the other and complete the exchange. However, leaving a harvester meant that any other village member would be able to claim it. This use of the spatial environment provides a natural implementation of a key friction in this environment: In order to engage in profitable exchanges, subjects may have to leave their harvesters and risk having their usage being coopted by others. This will have important implications for hypotheses concerning how patterns of production and exchange will be impacted under the 3H and 2H treatment conditions where the number of subjects per village exceeds the number of usable harvesters.

In order to facilitate interactions and exchange, subjects were given the ability to communicate with any avatar in a 20-meter radius using natural language. Another advantage of using our virtual environment is that it provided a set of contextual cues that could be used by subjects to coordinate with one another. Each individual avatar was given a generic appearance, and could be visually differentiated because the color of their shirt would match the color of their village. Furthermore, when avatars came sufficiently close to one another (within 50 meters) they would be able to see each other's names as tags floating about their heads, allowing for them to identify each other at a distance by shirt color and in proximity by user name. Subjects could use the cues provided in the environment to do things such as navigate through the virtual space, coordinate on meeting locations and direct one another towards harvesters.

In addition to these particular uses of the virtual environment to implement specific features of our design, the use of naturalistic framing in general can have important effects on subject behavior (Al-Ubaydli, McCabe and Twieg, 2014; Charness, Haruvy and Sonsino, 2007; Hoffman, McCabe and Smith, 1996; Kimbrough and Wilson, 2011). Additionally, Fiedler, Haruvy and Li (2011) and Atlas and Putterman (2011) specifically address the impacts of framing and social distance on subject behavior within Second Life, a virtual platform which is closely related to OpenSim.

4.1.4 Experimental Procedures

Actual experimental sessions were divided into six rounds, each lasting for 15 minutes. During these rounds, subjects would be able to engage in production and exchange. Subjects were not randomized to new avatars or villages between rounds, allowing for reputation and trust to be generated across the entire 90-minute session. At the end of each round, each subject's inventory was reset and whatever specializations they had acquired would have to be reacquired.

During the first two rounds of the experiment subjects could only engage in intra-village interactions and were not allowed to travel out of their home village. This meant that diamonds could not be produced during these rounds. From the third round onwards, subjects were allowed to leave their villages and engage in interactions with subjects from other villages. Subjects could reach these other villages either by traveling through openings in the walls between the villages (one of which can be seen in Figure 4.2) or by using teleporters in the center of each village that would provide transportation to the centers of the other two villages in a way that economized on the travel times related to inter-village journeys.

4.2 Measures and Hypotheses

The primary focus of this project is investigating how modulating an excess supply of labor under different treatment conditions affects the characteristics of both production and exchange that emerge within groups. Due to the absence of secure property rights over the resource-generating harvesters, subjects face a variable opportunity cost in switching from the production of one type of resource to another or in finding trading partners. A given subject i may not be able to utilize a particular production harvester in the future if the harvester used by i is claimed in her absence. The hypotheses outlined in this section will describe this experiment's key measures of interest and how modulating this scarcity is expected to affect these measures. Some of these measures concern overall group performance, but many of them attempt to examine how exactly these performance measures were impacted through the differing networks of trust and trade that emerge under different treatment conditions.

4.2.1 Income and Efficiency Measures and Hypotheses

The most important village performance measure that will be used in this analysis is *normalized* income. Because villages with additional harvesters have a higher potential for production, it would not be very illuminating to try to compare village performance across treatments by summing individual subject incomes - this result would be confounded by the potential income differentials that exist simply by virtue of having access to additional harvesters. In order to control for this confound, under each treatment condition I define a measure of expected *potential* income in order to provide a baseline of performance for that treatment. These expected potential income figures are derived from simulated experimental rounds that assume that subjects produce circles and triangles for 14 of the 15 minutes in each round, and then use the remaining minute to use these primary resources to create as many triangles as possible through exchange.[3]

Figure 4.4 depicts the distribution of simulated potential income using 5000 observations from each of the three treatment conditions. Under the 2H condition, simulated groups earned a total 7182 experimental dollars on average, with a standard deviation of 826 experimental dollars. Under the 3H condition, simulated groups earned a mean of 11831 experimental dollars on average, with a standard deviation of 991 experimental dollars. Under the 4H condition, simulated groups earned a total of 14379 experimental dollars on average with a standard deviation of 1179 experimental dollars. These averages are utilized in order to create the normalized income measures that will be used in this experiment's hypotheses - actual village income is divided by these potential income measures in order

[3]Note that in this idealized scenario there is no diamond production. However, if diamond production were possible this would just rescale all potential income metrics in a fashion that would have no qualitative impact for cross-treatment comparisons.

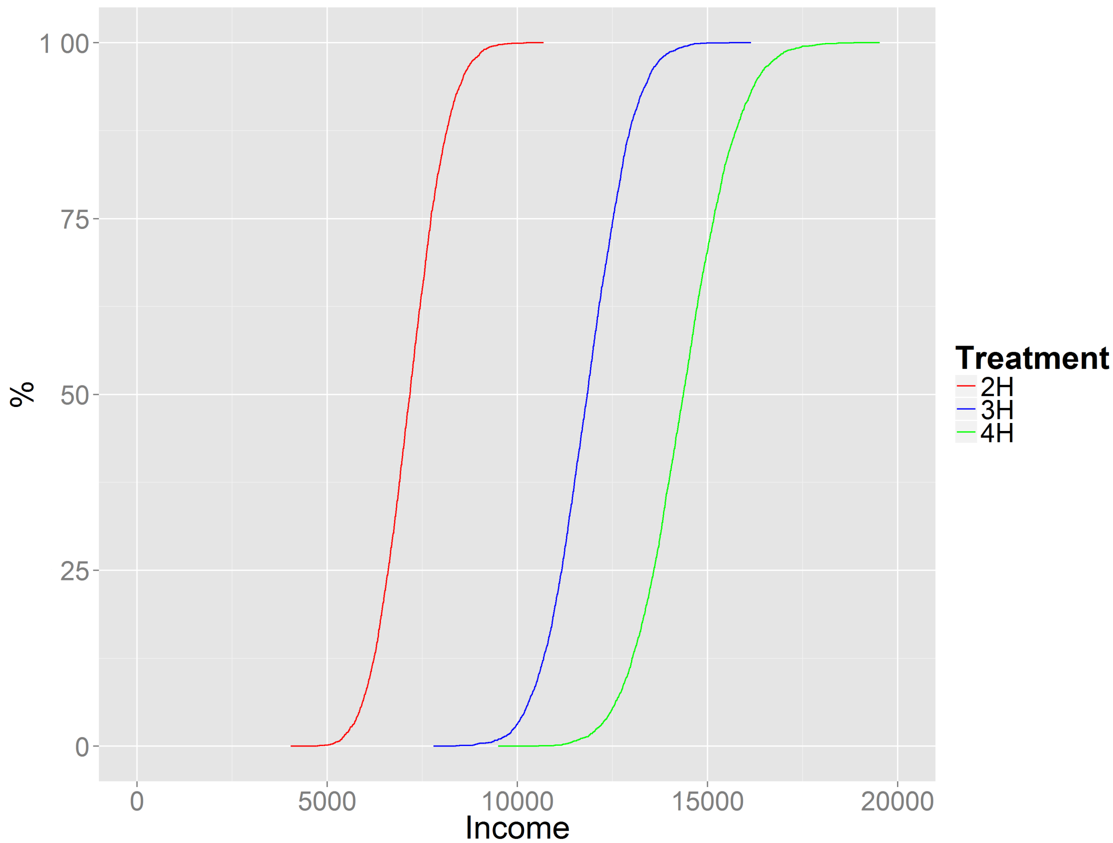

Figure 4.4: Cumulative distribution of simulated potential income with 5000 trials per treatment.

to generate normalized income.

With this in mind, this experiment's first hypothesis utilizes this normalized income measure:

Hypothesis 3. *Normalized income will fall monotonically with increasing labor supply.*

This hypothesis will be affirmed if subjects engage in suboptimal levels of exchange in the 2H and 3H treatment conditions due to concerns about losing access to production opportunities. However, one can also envisage institutions arising under which normalized incomes are higher in treatments where excess labor is present. This can occur when these excess village members serve not as threats to harvester-users' productive opportunities, but as effective trade intermediaries that allow for both high levels of production and exchange through a specialization of roles. In the 2H and 3H conditions some subjects could hypothetically spend the entire round producing at harvesters while having goods relayed to one another by non-producing subjects, assuming sufficient levels of trust and coordination exist within the group.

In this experiment, normalized village income can be understood both a product of both a village's ability to produce a large quantity of rings and circles, and an ability to effectively convert those rings and circles into triangles and eventually diamonds. The former capacity will be called *productive efficiency* and the latter will be called *exchange efficiency*. As an application of these concepts, the previous paragraph provided a narrative that hypothesized that excess labor would have a deleterious impact on exchange efficiency in particular - in order to test this dynamic, then, some additional definitions will be outlined and operationalized.

The main measure of productive efficiency that will be used is the total amount of primary resources (circles and rings) each village managed to generate per harvester. This metric captures multiple aspects of production at once: First, if harvesters are simply utilized for a larger portion of the total time in a round, more resources will be collected and this metric will be higher. Second, if subjects using the harvesters become more specialized, than total production will also be higher by virtue of the greater productivity that this specialization enables. It is with this definition in mind that this study's second hypothesis is outlined:

Hypothesis 4. *Productive efficiency will rise monotonically with increasing excess labor supply.*

Simply put, if subjects are unable to find exchange partners without giving up control of their harvesters, this may lead to their spending a larger share of the round at these harvesters, which will lead to greater production and specialization within groups. However, this increase in productivity will be offset to some degree by a decrease in exchange efficiency. In this analysis two broad measures are defined related to exchange efficiency: "Triangle efficiency" and "diamond efficiency." Triangle efficiency represents the share of possible triangles that could have been produced given a group's level of ring and circle production excluding the triangles that could have been produced in an autarkic fashion, while diamond efficiency will represent the share of possible diamonds that could have been produced given a group's level of ring and circle production as well as the circle and ring production of other groups.[4] Since diamond efficiency concerns intervillage trade, it is only a relevant measure for rounds 3-6 of a session. These measures aim to reflect how well groups actually realized their potential income given a fixed level of production. They also serve as a measure of deadweight loss. Note that in the simulated rounds used to generate the potential income measures for each treatment, triangle efficiency is always 100% after trading is completed.

Hypothesis 5. *Exchange efficiency (both triangle and diamond) will fall monotonically with increasing excess labor supply.*

If subjects are reluctant to leave their harvesters exchange efficiency will likely be negatively impacted. Even if the excess labor in the 2H and 3H treatment conditions were utilized to facilitate exchanges, this would not necessarily lead to higher exchange efficiency if trust and trustworthiness was lacking among these intermediaries or if there were high transaction costs involved with having a less direct flow of goods.

The next set of hypotheses will be concerned with exploring measures related to exchange efficiency and some of the institutions that groups may establish that will impact it. I will first be concerned with the extent to which the individuals within each village acquired triangles through either the use of multiple types of harvesters or through exchange. I define the *propensity to share* within a village as the share of the total amount of produced triangles that can be attributed to individuals using multiple harvesters and producing triangles in an autarkic fashion. In general, groups with a high propensity to share can be seen as engaging in an exchange of *harvesters* rather than of actual goods in a way that

[4]For example, if subject i produces 15 rings and 5 circles, and subject j produces 5 rings and 15 circles, then the total amount of triangles that could have been produced within this economy is 20. 10 of these could be produced via autarky, so triangle efficiency is 0% if subjects i and j jointly produce only 10 triangles. It is 100% if the full 20 are produced.

requires less bilateral trust in order to coordinate. However, this sort of exchange can be negatively impacted by a lack of property rights and an excess labor supply through the same mechanisms that underlie Hypothesis 5.

Hypothesis 6. *The propensity to share within groups will fall monotonically with increasing excess labor supply.*

When excess labor is present, the coordinated sharing of harvesters may become very difficult to engage in: If a harvester becomes available because a subject intends to try to use a different harvester, then there is both no guarantee that the subject that left will be able to use the other harvester nor is there a guarantee that they will continue to be able to reclaim the use of the old harvester if the other one can be attained. These factors should lower the propensity to share harvesters in the 2H and 3H treatment conditions.

Hypothesis 7. *The abandonment rate of harvesters will fall monotonically with increasing excess labor supply.*

Hypothesis 7 helps examine the asserted mechanism underlying Hypothesis 6. If the propensity to share is lower in the 2H and 3H treatment conditions because subjects are not using both harvester types, then we should less abandonment of individual harvesters within rounds. Abandonment herein is defined as a subject leaving a harvester for whatever reasons and thus allowing some other subject to claim it. If subjects are concerned that their ability to use harvesters in the future will be compromised by leaving to engage in some alternative activity, then this will result in a diminishment of the observed abandonment rate.

4.2.2 Trade Network Measures and Hypotheses

In order to further examine the characteristics of the exchange networks that emerge within and between villages, several additional measures and hypotheses will be defined. In order to aid with the formalization of these networks, let us define a set of transferable goods that exist in this environment as $R = \{$*Red Ring, Red Circle, Red Triangle, Green Ring, Green*

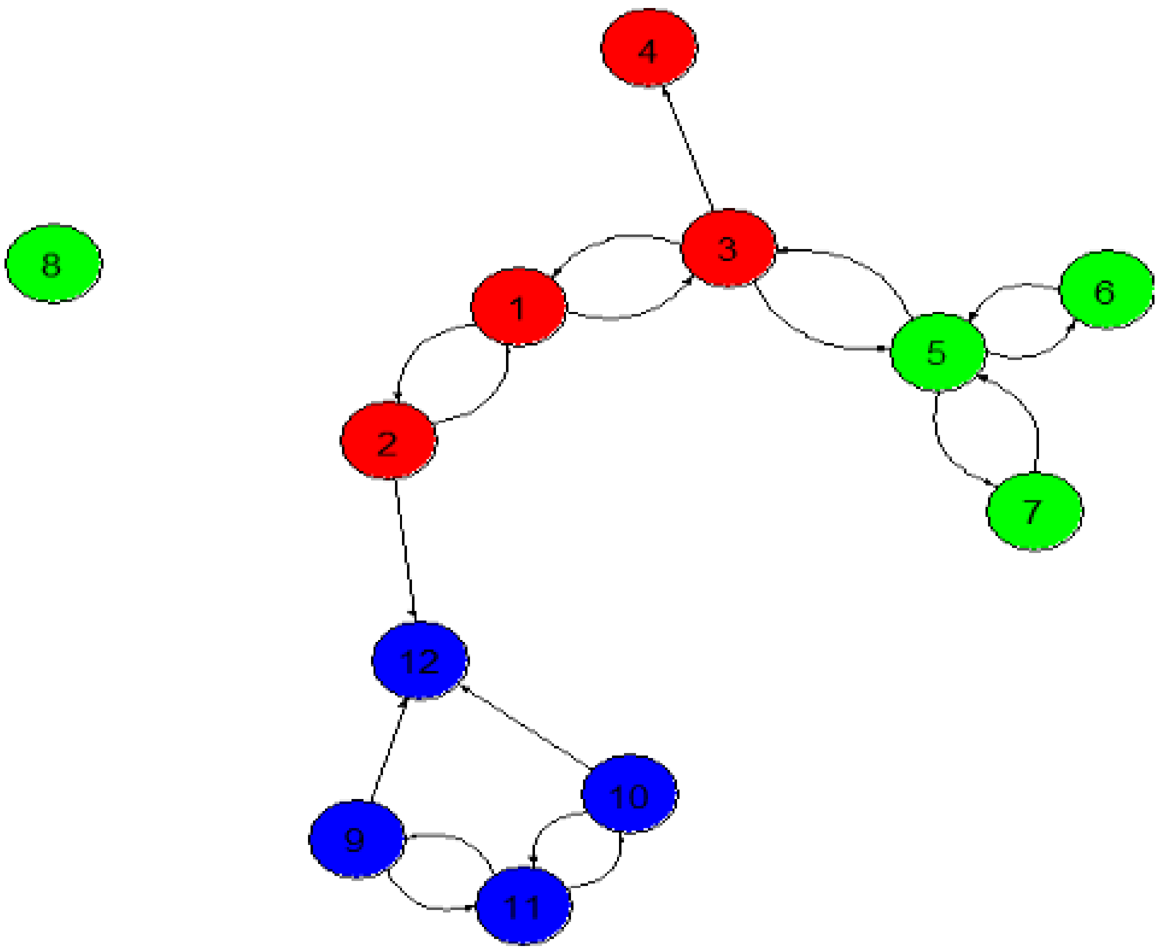

Figure 4.5: An example of an unweighted directed exchange network from a single experimental session-round.

Circle, Green Triangle, Blue Ring, Blue Circle, Blue Triangle, Diamond, Experimental Dollars}, and let $T^{r \in R}$ represent a 12x12 matrix wherein T^r_{mn} encodes how much of a good r was transferred from each subject m to subject n. T, in turn, represents the set of these 11 individual transfer matrices.

In general, the network statistics that will be discussed will be operationalized using a graph built on some $G(V, E(T))$. In this framework, V consists of a vertex for each subject within the experiment, and $E(T)$ is a mapping from T to a new 12x12 adjacency matrix by E. For example, Figure 4.5 shows an unweighted directed graph derived from a single round from this experiment that utilizes an adjacency matrix whereby an edge is drawn from subject i to subject j if subject i transferred any quantity of any good to j. Let the graph formed by this particular transformation be called $G'(V, E'(T))$, whose particular mapping $E'(T)$ can be formalized as:

$$E' = \min\left(\sum_{r \in R} T^r, 1\right) \tag{4.2}$$

The next two hypotheses concerning network formation under the different treatment conditions will utilize this network definition G':

Hypothesis 8. *The number of edges in G' involving each village's members will fall monotonically with increasing excess labor supply.*

Hypothesis 9. *The length of paths in G' involving each village's members will rise monotonically with increasing excess labor supply.*

If exchange efficiency falls under the 2H and 3H treatment conditions, then this will likely correspond to a degradation in the observed exchange networks represented by G': There will be fewer trade links (Hypothesis 8) and the path lengths between any agents i and j that are connected by a trade link will become more indirect (Hypothesis 9.) Note that despite G' being an intervillage network, these hypotheses concern measures that are done at the village level. These measures are generated by focusing in on all of the nodes associated with individual villages within the network. For example, rather than looking at the path lengths between each possible i and j, the village-level path length measure looks at all paths between i, j pairs where either i or j belong to the given village.

The next three hypotheses represent refinements of the previous two that investigate the nature of these trade paths in greater detail. A key question for this investigation is whether the presence of excess labor leads to the emergence of intermediaries in trade. In order to test any hypothesis concerning intermediaries, some additional definitions will now be offered. Let us call a subject i a *circle intermediary* if i transferred a larger quantity of circles of their village color than they acquired through harvesters. Similarly, let us call a subject i a *ring intermediary* if the same condition applied for rings. In practical terms, a subject can only qualify as a circle or ring intermediary if they both sent out the relevant good and were given some quantity of the good by others, thus acting as a conduit through which those goods were transferred on to a third party. If subjects were truly acting as intermediaries between a party with an excess supply of rings and a party with an excess supply or circles, then one would expect them to be categorized as both circle intermediaries and ring intermediaries within a single round subject i would take rings from subject j to subject k, and would then take circles back from k to j. Going forward, subjects will be referred to as *both intermediaries* if they simultaneously fulfill both the ring intermediary and circle intermediary criteria. Using these definitions the next hypothesis can be established:

Hypothesis 10. *The likelihood with which any intermediaries will be observed will rises monotonically with increasing excess labor supply.*

The last two hypotheses concern how subjects who engage largely in production are connected to one another under the different treatments. To formalize the necessary measures to test these hypotheses, each subject is assigned to a role given their degree of specialization:

- A subject is defined as a *ring producer* if they achieved at least one level of specialization in rings but no specialization in circles. Similarly, a subject is defined as a *circle producer* if they achieved at least one level of specialization in circles but no specialization in rings.
- If a subject achieved a level of specialization in both circles and rings, they are categorized as a *both producer*.
- If a subject achieved no specialization, they are categorized as a *neither producer*.

Note that this typography allows for subjects to be classified as both ring/circle producers and intermediaries, although one might not expect for this to occur often. In terms of these productive types, then, I investigate two additional hypotheses:

Hypothesis 11. *Trade path lengths between subjects assigned the circle producer role and subjects assigned the ring producer role within the same village will rise monotonically with increasing excess labor supply.*

Hypothesis 12. *The likelihood with which a ring producer and a circle producer will be connected by a trade path within a village will fall monotonically with increasing excess labor supply.*

These hypotheses investigate whether the producers of different primary goods are more distant from one another when excess labor is present, if they are connected at all. The goal of Hypotheses 11 and 12 is to more-precisely investigate how the labor supply surplus that exists in the 2H and 3H treatment conditions affects whether *certain* key vertices are connected G', and how. Collectively, the goal of the hypotheses outlined in this section is to facilitate an exploration of not only how overall village welfare is impacted by excess labor supply but how this excess labor supply qualitatively impacts the emergent characteristics of production, specialization, and exchange in the experimental environment.

4.3 Statistical Design and Data Measurement

The data for this analysis comes from fifteen experimental sessions that took place in the George Mason University Krasnow Institute laboratory from the Spring 2011 to the Spring 2012 semesters. Each session lasted for two and a half hours in total, beginning with an hour of software training followed by six 15-minute rounds of the experimental procedure. Subject instructions can be found in Appendix C. Each session included twelve subjects in total who were divided into three villages of four. In the first two rounds of the experiment, subjects could only interact with members of their own villages and produce rings, circles, and triangles. In the third round and beyond, subjects could interact with members of the other villages and produce diamonds. All villages in a session experienced the same treatment condition: 2H, 3H, or 4H.

The fundamental unit of observation in this investigation is the village, and all hypothesis tests in this study will be performed using village-level observations. As a result, the dataset for this experiment contains 270 village-round observations in total, 90 from each of the three treatment conditions. With a couple minor exceptions, the subsequent tests all use the following econometric specification to estimate the treatment effects that will serve to accept or reject the outlined hypotheses:

$$y_{it} = a + b^{2H}T_i^{2H} + b^{3H}T_i^{3H} + \sum_{s=2}^{6} b_T^s T_t^s + \mu_{g(it)} + \mu_{s(it)} + \epsilon_{it} \tag{4.3}$$

Where:

- y_{it} is the relevant dependent variable of interest for village i in round t.
- a is the intercept term.
- T_i^{2H} and T_i^{3H} are indicator variables for the treatment condition experienced by village i. $4H$ is the baseline treatment condition.
- T_t^s represents a vector of period indicator variables for periods 2 to 6.
- $\mu_{g(it)}$ is a village-level cluster variable for i which addresses village-specific autocorrelation across the 6 round observations derived per village.
- $\mu_{s(it)}$ is a session-level cluster variable for i which addresses session-specific autocorrelations across the 4 observations per village where the different villages in a session could interact with one another.

Note that because the previously-outlined hypotheses pertained to the monotonicity of different dependent variables with respect to increasing excess labor supply, there are essentially three specific hypothesis tests that will be of interest to discussing monotonicity, namely:

1. $b^{2H} = 0$ - This null hypothesis pertains to the 2H vs. 4H treatment condition.

2. $b^{3H} = 0$ - This null hypothesis pertains to the 3H vs. 4H treatment condition.

3. $b^{3H} - b^{2H} = 0$ - This null hypothesis pertains to the 2H vs. 3H treatment condition.

The next section will discuss all three of these tests when estimating treatment effects on variables of interest. However, for the sake of simplicity all the tables in the next section show coefficients and standard errors using the 4H treatment as a baseline - this allows for the easy comparison of the 4H treatment condition to the 2H and 3H treatment conditions. The 2H and 3H conditions can additionally be compared by dividing the difference between their estimated coefficients by their common standard error. In general merely affirming (1) will be taken to demonstrate monotinicity assuming that $b^{3H} - b^{2H}$ is not both statistically significant and opposite in sign from b^{2H}. This is a fairly weak standard of monotinicity, and it will be discussed when stronger definitions can be satisfied given the observed treatment effects.

4.4 Data Analysis

4.4.1 Income and Efficiency Results

Result 1. *Normalized income within villages does not fall monotonically with increasing excess labor supply.*

This result is shown in Table 4.1 as well as Figure 4.6. Recall that these normalized incomes represent a proportion of the potential income figures discussed earlier. The mean normalized income across all treatments is roughly .4, and the only marginally significant treatment effect that exists on this metric indicates that the 2H treatment condition achieves a higher share of normalized income than the 3H treatment condition ($p < .1$). It is particularly noteworthy that many of the highest-performing groups on this metric belonged to the 2H condition - of all 7 village-round observations that achieved a normalized income measure of .7, all of them occurred in 2H villages.

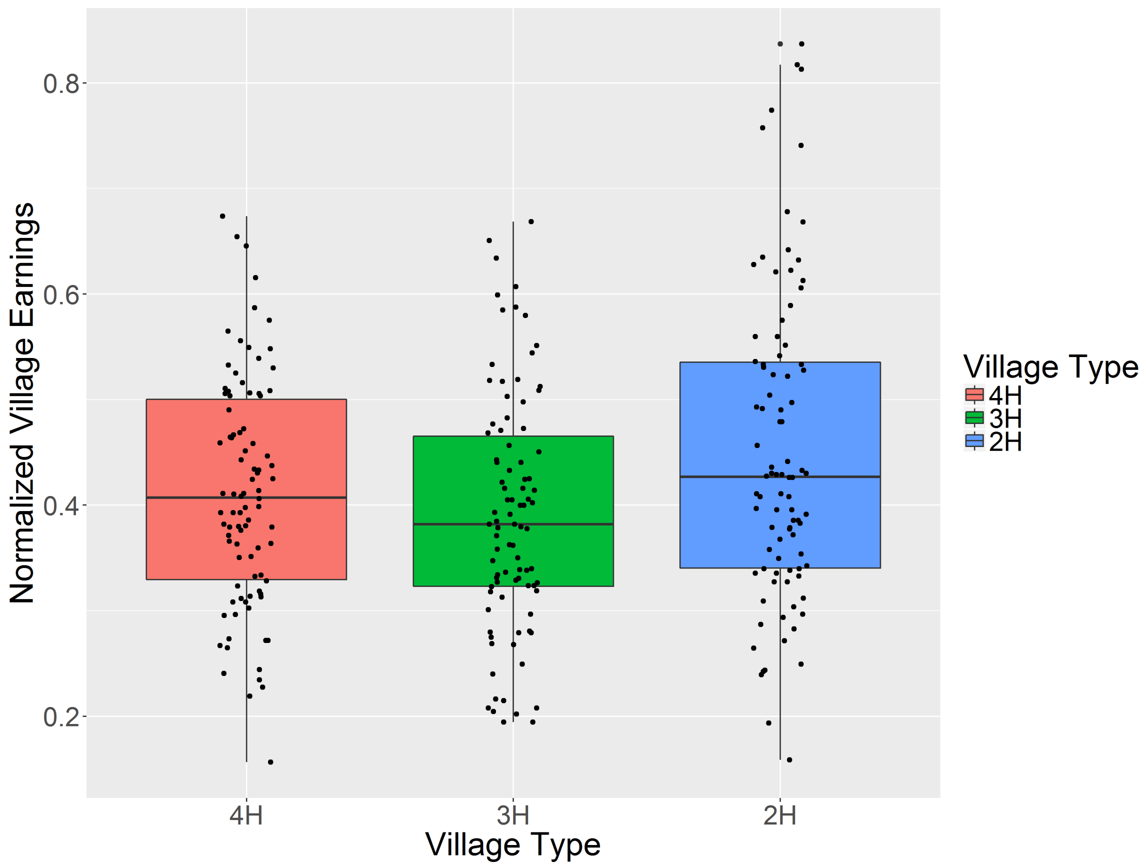

Figure 4.6: Distribution of normalized village income by treatment.

These finding means that the null hypothesis corresponding to Hypothesis 3 cannot be rejected - rather than observing that an excess supply of labor resulted in a decrease in normalized incomes, the statistical evidence indicates that this metric may increase with excess labor. Several of the subsequent results will provide insight into why exactly this is occurring.

Result 2. *Productive efficiency rises monotonically with increasing excess labor supply.*

This result is shown in Table 4.2 as well as Figure 4.7. Table 4.2 includes both the overall productive efficiency measure and the productive efficiency measures that would be generated if only rings or only circles were examined. This effect is mainly evidenct when comparing the 4H treatment condition to the 2H or 3H treatment condition - the

Table 4.1: Estimated effects on normalized village income.

	Normalized Village Income
Intercept	0.322^{***}
	(0.024)
3H Treatment	-0.022
	(0.031)
2H Treatment	0.037
	(0.031)
Round 2	0.080^{***}
	(0.018)
Round 3	0.065^{***}
	(0.018)
Round 4	0.096^{***}
	(0.018)
Round 5	0.129^{***}
	(0.018)
Round 6	0.176^{***}
	(0.018)
Num. obs.	270

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

difference between 2H and 3H is not statistically significant. Overall, the 2H treatment condition produces over 30% more circles and rings per harvester than the 4H condition. One additional result is that productive efficiency did not increase substantially across the 6 rounds of the experiment - it's only in the last couple rounds that a marginally significant uptick in production can be observed.

This result raises a natural question: If productive efficiency increased under the 2H and 3H treatment conditions relative to the 4H, why did this not translate into an observable impact on normalized income? As the next set of results show, this is likely due to the fact that there was a compensating negative impact on the exchange networks that emerged under those treatments.

Result 3. *Exchange efficiency (both in triangles and diamonds) falls monotonically with increasing excess labor supply.*

This result is shown in Table 4.3 and visualized in Figures 4.8 and 4.9. Once again

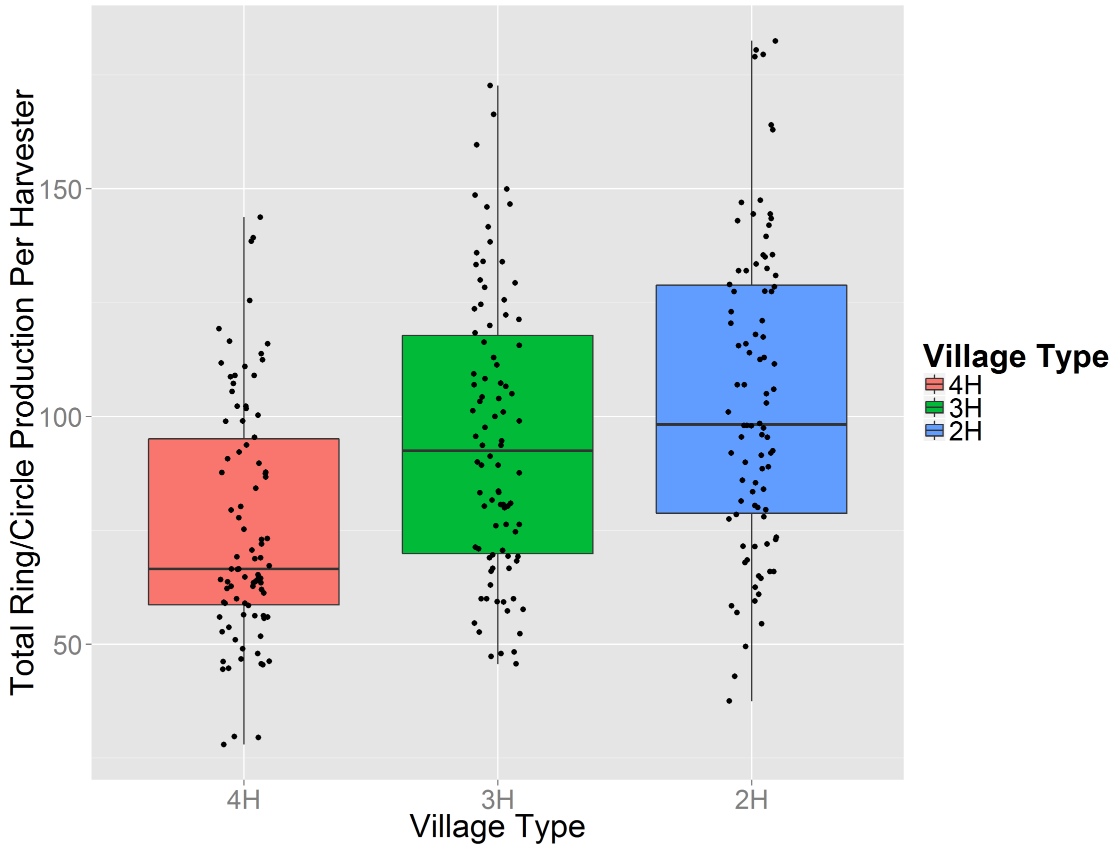

Figure 4.7: Productive efficiency by treatment.

the 2H and 3H conditions are statistically identical, but villages under the 4H condition were substantially more effective at converting excess rings and circles into triangles and diamonds via exchange. Figure 4.8 corroborates this point by showing the distributions of triangle efficiency scores across the three treatments - of the 24 village-rounds that achieved perfect triangle efficiency, over half of them occurred under the 4H treatment condition. In general diamond efficiency was substantially lower than diamond efficiency, likely reflecting the added difficulty of having to extend trade paths even further to acquire diamonds. As seen in Figure 4.9 only 16 village-rounds achieved 30% efficiency in acquiring diamonds, and of those 16 observations 11 of them came from the 4H treatment condition.

Result 4. *The propensity to share within groups falls monotonically with increasing excess*

Table 4.2: Estimated effects on productive efficiency measures.

	Average Harvester Production	Average Ring Harvester Production	Average Circle Harvester Production
Intercept	71.343^{***}	55.478^{***}	79.648^{***}
	(6.746)	(5.277)	(11.079)
3H Treatment	16.912^{*}	12.483^{*}	50.728^{***}
	(8.659)	(6.468)	(13.684)
2H Treatment	25.366^{***}	17.183^{***}	37.461^{***}
	(8.659)	(6.468)	(13.684)
Round 2	16.120^{***}	9.744^{**}	25.956^{***}
	(4.508)	(4.187)	(8.358)
Round 3	1.154	−1.556	9.956
	(4.998)	(4.616)	(8.358)
Round 4	2.700	0.689	11.167
	(4.998)	(4.616)	(8.358)
Round 5	9.052^{*}	5.567	18.489^{**}
	(4.998)	(4.616)	(8.358)
Round 6	9.024^{*}	5.622	18.111^{**}
	(4.998)	(4.616)	(8.358)
Num. obs.	270	270	270

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

labor supply.

Result 5. *The abandonment rate of harvesters falls monotonically with increasing excess labor supply.*

These two results are shown in Table 4.4. In the 4H treatment condition, slightly over half of the triangles that could have been produced within villages can be attributed to subjects sharing the harvesters. These are triangles that required no exchange at all in order to be acquired within the village. However, in the 2H and 3H treatment conditions this figure is just below 30%, indicating that the presence of excess labor substantially diminishes the extent of harvester sharing within these villages. Furthermore, harvester sharing also exhibits a relative decline in its importance in the latter rounds of the experiment, particularly in round 4 and onwards as groups develop reliable exchange networks.

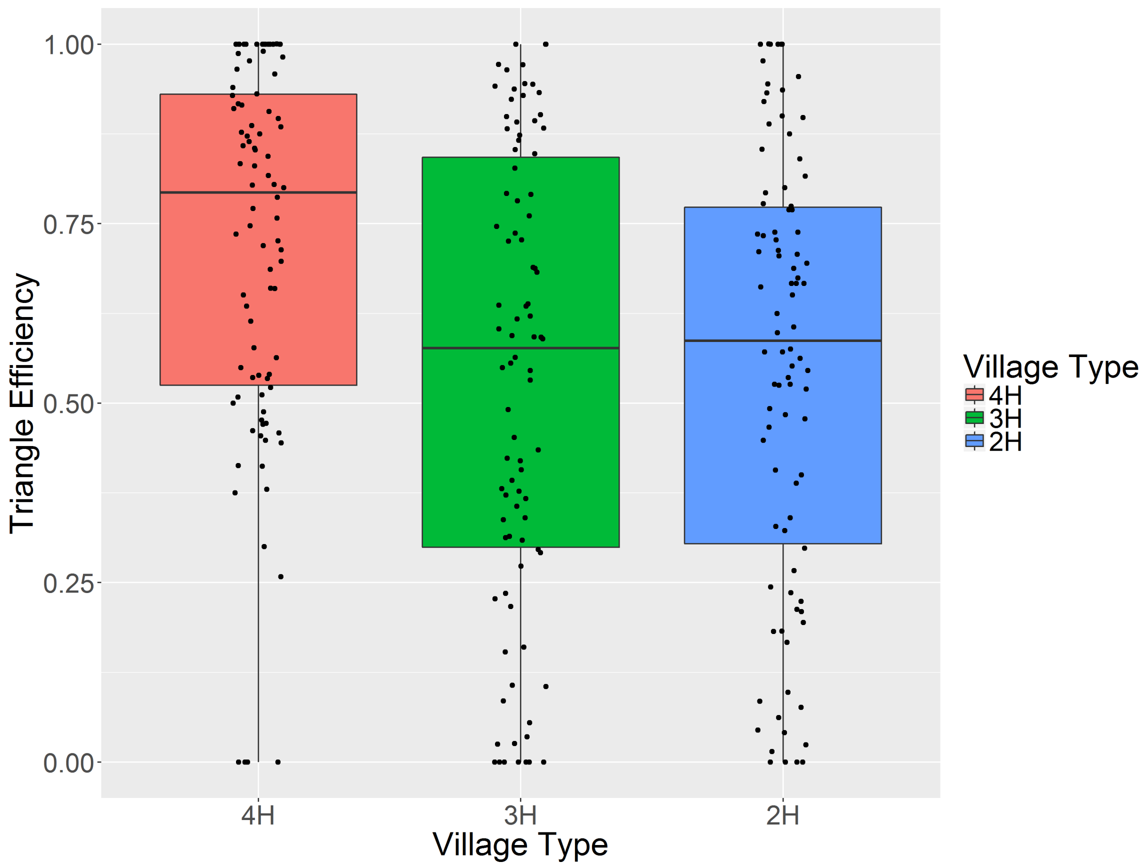

Figure 4.8: Triangle exchange efficiency by treatment.

This increased sharing of harvesters also corresponds to an increase in the average number of times that harvesters are abandoned in a round. There is strong monotinicity in the abandonment rate across different treatments, with the 2H condition exhibiting significantly less abandonment than the 3H condition, which in turns exhibits less abandonment than the 4H treatment condition. This indicates that the presence of excess labor under these conditions is causing subjects to become more cautious about leaving harvesters and giving up their tenuous squatters' rights over them.

Taken together, this first set of results indicate that the 2H and 3H treatment conditions caused villages to engage in more production but less exchange, despite the excess labor supply: Total per-harvester production of primary resources was higher in the presence of an excess labor supply, but fewer of these resources were converted into higher-order goods

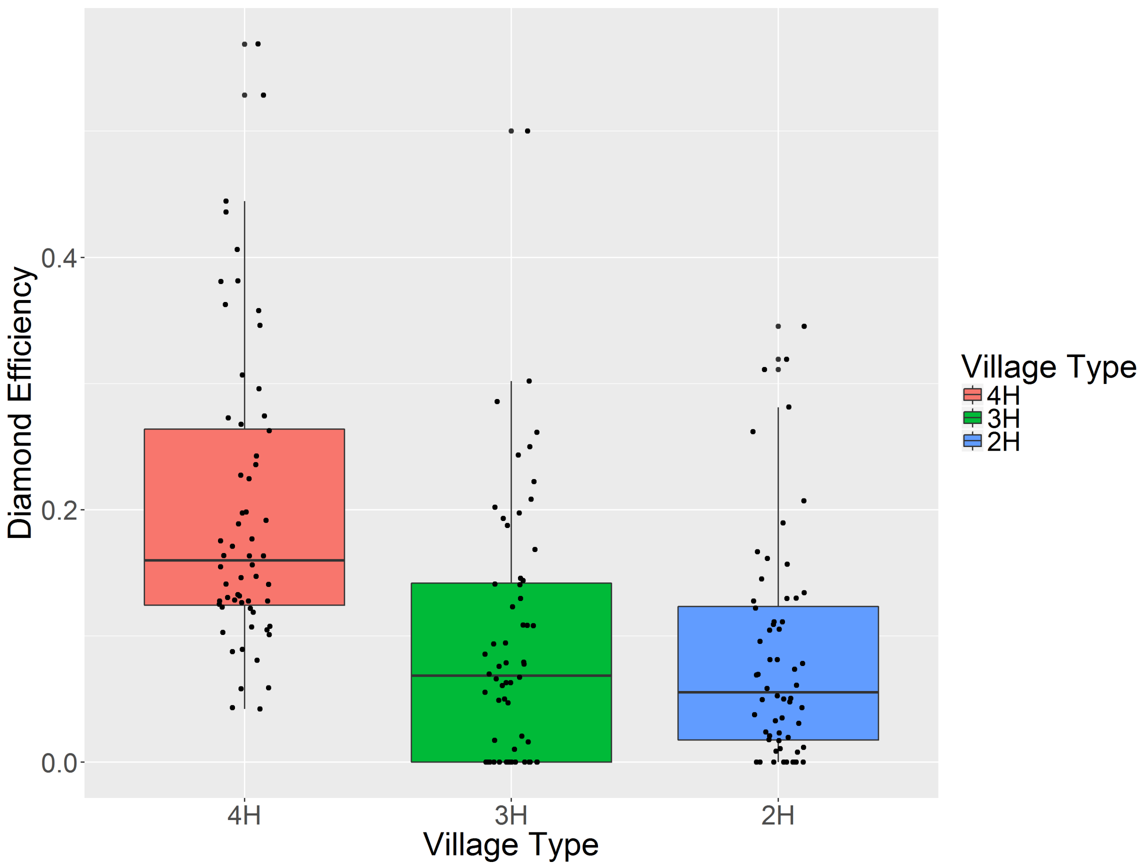

Figure 4.9: Diamond exchange efficiency by treatment.

via exchange. The remaining results will describe how these treatment conditions affected the nature of the trade networks that emerged as well as how subjects adopted various productive and intermediary roles. These results will provide insights into why exactly this important drop in exchange efficiency was observed.

4.4.2 Trade Network Results

Result 6. *The number of directed edges in G' involving each village's members does not fall monotonically with increasing excess labor supply.*

This result is shown in Table 4.5. The first column of this table represents estimates of the number of directed edges of a given unweighted exchange network G' among a given

Table 4.3: Estimated effects on exchange efficiency measures.

	Triangle Exchange Efficiency	Diamond Exchange Efficiency
Intercept	0.626^{***}	0.041^{**}
	(0.060)	(0.017)
3H Treatment	-0.179^{**}	-0.058^{***}
	(0.071)	(0.020)
2H Treatment	-0.161^{**}	-0.064^{***}
	(0.071)	(0.020)
Round 2	0.100^{*}	-0.000
	(0.052)	(0.016)
Round 3	0.036	0.108^{***}
	(0.056)	(0.018)
Round 4	0.117^{**}	0.110^{***}
	(0.056)	(0.018)
Round 5	0.079	0.127^{***}
	(0.056)	(0.018)
Round 6	0.193^{***}	0.157^{***}
	(0.056)	(0.018)
Num. obs.	270	270

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

session-round that pertain to a villagers belonging to individual villages. Unsurprisingly, this regression shows that the number of edges grows substantially (and continues to grow) after the second round, reflecting the development of intervillage trade. The second column estimates the number of edges in G' that connect members of the same village and shows more inconsistent growth across time.

It might be surprising that the number of edges in G' does not exhibit any treatment effects. The implications of this finding merits further discussion. One possible interpretation of a rejection of Hypothesis 8 is to assert that the presence of excess labor may not lead to a significant reduction of total transfer pairs as measured in G'; however, each individual edge of a trade network may still be less effective in generating wealth through combinations when excess labor is present. Imagine a scenario A where subject i produces 20 rings and subject j produces 20 circles. Subjects k and l do nothing. If i trades 10 rings

Table 4.4: Estimated effects on production sharing measures.

	Share of Triangles Produced by Sharing	Harvester Abandonment Rate
Intercept	0.562^{***}	3.014^{***}
	(0.078)	(0.204)
3H Treatment	-0.263^{***}	-0.748^{***}
	(0.098)	(0.177)
2H Treatment	-0.266^{***}	-1.140^{***}
	(0.098)	(0.177)
Round 2	-0.146^{***}	-0.937^{***}
	(0.055)	(0.184)
Round 3	-0.058	-1.096^{***}
	(0.056)	(0.184)
Round 4	-0.101^{*}	-1.287^{***}
	(0.056)	(0.184)
Round 5	-0.157^{***}	-1.266^{***}
	(0.056)	(0.184)
Round 6	-0.165^{***}	-1.300^{***}
	(0.056)	(0.184)
Num. obs.	270	270

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

to j in exchange for 10 circles, then this will be measured as a village network with a single edge.

However, now imagine a scenario B where k serves as an intermediary between i and j. i gives their 10 rings to k, who takes them and passes some them to j, who in return gives 10 circles to k to return to i. Under this scenario, one would observe a network with 2 edges. This edge measure may give the impression that more intravillage trade exists in scenario B than in scenario A, but in an important sense this is not the case the presence of the intermediary k leads to double-counting of each resource unit's journey through the trade network. The path measures that are discussed next aim to investigate these concerns. For now it should be observed that the reason for lower exchange efficiency in the 2H and 3H treatment conditions does not appear to be attributable to subjects simply being unable or unwilling to engage in any sort of exchange at all.

Table 4.5: Estimated effects on directed trade network edge measures.

	All Trade Edges	Intravillage Trade Edges
Intercept	4.691^{***}	2.004^{***}
	(0.678)	(0.266)
3H Treatment	-1.052	0.155
	(0.834)	(0.640)
2H Treatment	0.049	0.889
	(0.834)	(0.640)
Round 2	1.422^{***}	1.422^{***}
	(0.549)	(0.433)
Round 3	3.911^{***}	-0.044
	(0.549)	(0.433)
Round 4	4.622^{***}	1.066^{**}
	(0.549)	(0.433)
Round 5	5.021^{***}	0.578
	(0.549)	(0.433)
Round 6	6.043^{***}	1.469^{***}
	(0.549)	(0.433)
Num. obs.	270	270

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Result 7. *The length of trade paths in G' involving each village's members rise monotonically with increasing excess labor supply.*

This result is shown in Table 4.6, which again looks at both intravillage-only network links and all network links. Although trade paths tend to be fairly short between agents, they are longer in the 3H and 2H treatment conditions, indicating that the presence of excess labor is causing exchange to become indirect. Although intravillage trade paths do not exhibit a tendency to become longer throughout sessions, intervillage paths experience a significant increase in length between rounds 3 and 6. Note that these measures only have values when any valid links exist - the small number of observations where no exchange whatsoever occurred were dropped from this analysis. The distribution of missing values in this analysis did not exhibit any correlation with their treatment condition.

Result 8. *The likelihood with which any intermediaries are observed rises with increasing excess labor supply.*

Table 4.6: Estimated effects on trade network path length measures.

	Intravillage Path Lengths	All Path Lengths
Intercept	1.029***	1.040***
	(0.029)	(0.030)
3H Treatment	0.079**	0.069**
	(0.036)	(0.034)
2H Treatment	0.109***	0.084**
	(0.036)	(0.034)
Round 2	0.039	0.018
	(0.037)	(0.036)
Round 3	0.082**	0.194***
	(0.037)	(0.036)
Round 4	0.020	0.223***
	(0.037)	(0.036)
Round 5	0.034	0.236***
	(0.037)	(0.036)
Round 6	0.030	0.307***
	(0.037)	(0.036)
Num. obs.	254	259

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

This result is shown in Table 4.7, which uses a probit model to estimate the probability that any sort of intermediary (circle, ring, or both) will be observed in a given village-round. The 2H treatment condition exhibits a significantly-greater likelihood of having at least one intermediary: Of the 83 village-round observations where an intermediary is present, over half of them occur under the 2H treatment condition, indicating that the presence of excess labor is leading to the emergence of specialized intermediaries in the experimental environment.

Result 9. *Trade path lengths between subjects assigned the circle producer role and subjects assigned the ring producer role within the same village rise monotonically with increasing excess labor supply.*

Result 10. *The likelihood with which a ring producer and a circle producer will be connected by a trade path within a village falls monotonically with increasing excess labor supply.*

Table 4.7: Estimated effects on the probability of any intermediaries emerging.

	Intermediary Probability
Intercept	-1.558^{***}
	(0.351)
3H Treatment	0.501
	(0.343)
2H Treatment	0.989^{***}
	(0.345)
Round 2	0.902^{***}
	(0.325)
Round 3	0.077
	(0.343)
Round 4	0.179
	(0.339)
Round 5	0.772^{**}
	(0.327)
Round 6	0.528
	(0.331)
Num. obs.	270

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

These results are shown in 4.8. Like Table 4.7, there are a number of missing observations here - village-rounds without both a ring and a circle are excluded from the link probability estimate in the second column of this table, and observations without a trade path between such specialists are excluded from the path length estimates. Result 9 is shown in the first column of Table 4.8. Of all the links between between ring and circle specialists, over 80% of them are direct (having a length of one.) However, this figure rises to 90% under the 4H treatment and 75% under the 3H and 2H treatments, which are statistically indistinguishable. This finding dovetails intuitively with Result 8, since longer trade paths indicate the likely presence of intermediaries.

The second column of Table 4.8 provides the basis for Result 10. Given the existence of a circle specialist and a ring specialist within a village, the probability of their being linked by an edge in G' is significantly lower in the 3H and 2H treatment conditions than the 4H treatment condition.

Table 4.8: Estimated effects on path length between specialists and the probability that these links will emerge.

	Specialist Path Length	Specialists Linked Probability
Intercept	0.955^{***}	0.844^{***}
	(0.086)	(0.077)
3H Treatment	0.192^{*}	-0.184^{**}
	(0.102)	(0.075)
2H Treatment	0.281^{***}	-0.188^{**}
	(0.101)	(0.075)
Round 2	0.113	-0.031
	(0.069)	(0.073)
Round 3	0.113	0.021
	(0.077)	(0.073)
Round 4	0.047	0.030
	(0.076)	(0.074)
Round 5	0.173^{**}	0.021
	(0.077)	(0.074)
Round6	0.117	-0.069
	(0.075)	(0.077)
Num. obs.	198	230

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

4.5 Discussion and Conclusion

The unifying narrative of these results indicates that although excess labor supply did not directly impact overall group income, it had the expected effect of lowering the effectiveness of exchange. Several of the measures used in this investigation were meant to provide observable proxies for whether the hypothesized pressures that would arise due to excess labor supply were present. For example, the finding that harvester abandonment exhibits a significantly monotonic relationship with excess labor supply is taken as strong evidence that subjects are responding to the elevated threat of losing access to productive opportunities should they leave. Likewise, the finding that the trade network G' has as much edges under each treatment condition but that these edges are less likely to link different types of production specialists lends evidence to the idea that exchange efficiently is lower in these villages because the trade networks that emerged did not build paths between the most important sets of vertices because these individuals could not exchange without abandoning their harvesters.

One general finding is that for most measures, if any treatment effect was present the presence of *any* excess labor was sufficient to generate it. The only metric where the 2H treatment condition had an effect but the 3H condition did not was on the intermediary probability measure. For the most part, strong monotinicity where each treatment condition was significantly distinct was not present. What appeared to be important was the presence of any excess labor whatsoever. Ultimately, by way of leveraging a virtual world platform to instantiate an environment where the tradeoffs between engaging in production and engaging in exchange are spatiotemporal in a naturalistic fashion and excess labor supply is subject to exogenous variation, these results serve to provide insight into the mechisms through which weak property rights can have a deleterious impact on the networks of truth and exchange that emerge between individual actors.

Appendix A: Mathematical Appendix for Chapter 2

1-Player Planner Strategy

Because the growth of all berry bushes are independent and M is a linear function of the yields of all the individual bushes and the planner is assumed to be able to harvest all the berry bushes in a single period, the solo planner's problem of optimally maintaining n bushes is identical to the problem of maintaining a single berry bush as the same strategy will be employed over all n bushes in each period. Looking at one berry bush allows for the state space to be reduced in complexity from 10^{48} separate states to only 10. Let s^t refer to the number of berries on this one berry bush in period t. It is obvious that $b^T = s^T$. The solo planner will clear all of the berries in the last period since there is no return to the planner of having berries growing on bushes at the end of the game. This leads to a payoff in period T of $r\rho s^T$.[1]

In period $T-1$, the planner will choose b^{T-1} to maximize the value function outlined in Equation A.1:

$$V^{T-1}(s^{T-1}) = \max_{b^{T-1} \leq s^{T-1}} \{r\rho b^{T-1} + \sum_{i=0}^{10} P(s^T = i | s^{T-1} - b^{T-1}) r\rho s^T\} \tag{A.1}$$

Intuitively, the marginal cost or benefit of harvesting in period $T-1$ is that it affects the expected amount of growth that will occur on the berry bush between $T-1$ and T. The planner maximizes revenue by choosing b^{T-1} to maximize this growth. As mentioned in the discussion of Equation 3.1, the growth-optimizing value is a function of the transition matrix $H(\phi)$. Staying in the one berry bush case, if $s^t = a$ then equation A.2 defines the expected marginal growth from period t to $t+1$:

[1] I am assuming here that the planner is specialized in berries of the color of bush k here. If this is not the case, the analysis is identical and I can just scale down the payoffs received under each strategy by r.

$$s^{'}(t) = -s^t + E(s^{t+1}) = -s_t^a + \sum_{j=1}^{10} iH_{aj} \tag{A.2}$$

The value that maximizes s' here also depends on ϕ. If ϕ is close to 0, then the relationship between current berries and the expected growth rate is similar to the one depicted in Figure 3.1 with a maximum at $a = 4$ and declining monotonically around this point. The value $\phi = 0.0025$ used in the experiment fulfills this criteria. However, if ϕ is close to 1, then s' will be maximized at a level lower than 4. Intuitively, ϕ being close to 1 describes the case where berry growth is so certain that merely having any berries left on a bush will lead to most of the berries growing back in the next period. Going forward it will be assumed that $\phi = 0.0025$. Equation A.3 dictates that in period $T-1$ the planner chooses $b^{T-1} \leq s^{T-1}$ in order to maximize the sum of the present value of the berry bush and the expected change in this value from harvesting b^{T-1} berries:

$$V^{T-1}(s^{T-1}) = r\rho b^{T-1} + \max_{b^{T-1} \leq s^{T-1}} \{s^{'}(T-1)\} \tag{A.3}$$

Since $r\rho b^{T-1}$ is constant and $s^{'}$ is quadratic in s^t, the sum of the two is also quadratic and $V^{T-1}(s^{T-1})$, like s', is maximized when b^{T-1} is chosen in order to set $s^{T-1} - b^{T-1}$ as close to 4 as possible. Two important things are of note here. Firstly, when $s^{T-1} > 4$, then $V^{T-1}(s^{T-1}) = V^{T-1}(4) + s^{T-1} - 4$ - that is, since the planner brings the number of berries on the bush down to 4 if it is above this amount, then entering period $T-1$ with more than 4 berries simply translates into extra berries being harvested in $T-1$. Through recursion, this also applies to any period prior to $T-1$: If in period t the planner knows that it is optimal to set s^{t+1} to 4, then if $s^t > 4$ the planner's sole concern will be to choose b^t in order to maximize $s^{'}(t)$. This will be achieved by harvesting the bush down to 4 berries. Secondly, for $s^{T-1} < 4$ it is the case that $V^{T-1}(s^{T-1})$ increases with s^{T-1}. If in any period t where $s^t < 4$ we know that V^{t+1} is increasing in s^{t+1}, then it will always be optimal to abstain from harvesting - not only will the unharvested berries be available in the next period, but not harvesting will also maximize $s'(t)$ and improve the planner's chances of achieving a more-valuable berry bush state in the next period. These observations jointly show that the optimal management rules in $T-1$ also hold in prior periods.

4-Player Planner Strategy

Having shown that the solo planner adopts the management rule $g(4,T)$, this result can be generalized to the 4-player case that is used in the experiment. Once again, payoffs are linear in berries collected if the specializations of each of the four subjects is the same as the specialization of the planner in the 1-player case. Any pattern of harvesting will yield the same total income in either the 1-player or the 4-player planner case. Therefore any combination of individual management strategies that collectively implements the 1-player planner strategy will represent an efficient outcome. Notably, this can include strategies where the subjects partition berry bushes and only implement the $g(4,T)$ rule on the berry bushes in their partition while leaving all the others alone. In order to resolve cases where two subjects may be making simultaneous choices to harvest the same berries, I implement a simple tie-breaking mechanism. I assume that in each period Nature randomly chooses an order through which each subject's harvesting decisions are implemented. If, for example, all four subjects follow the $g(4,T)$ rule, then it will be the case that the first subject chosen by Nature will end up harvesting all the berries in a single period.

In the no cost specialization treatment, each subject will choose to specialize in all colors other than silver. In the costly specialization treatment, assuming that at least one berry of each of the RGBY types is present, then maximum payoffs are achieved by having each subject choose a different one of these four colors and not harvest any berry bush that they are not specialized in. This represents a partitioning of the RGBY bushes based on type. Silver berry bushes can be partioned in any fashion. Under both treatments both the management rules and total payoffs will be identical in expectation.

Nash Equilibrium

Any strategy that a subject can employ that doesn't specify a clearing time τ will be dominated by one that does, since any strategy that clears the field at period T can be characterized as setting $\tau = T$ and not clearing the field at the end will be dominated by any strategy that does. Assume the case where subjects are not immediately clearing the field - where $\tau > 1$ for all subjects. Let τ_{min} represent the lowest value of τ from among the subjects, and let subject i represent one of the subjects who is using the cutoff time τ_{min}. No matter what strategies the other subjects employ, each other subject will expect to receive at best half of the field's value in period τ, assuming they adopt a clearing time of τ and accept at most a 50% chance of being able to clear the field before another subject.

Let Π_j^t represent the total value of the set of berry bushes to subject $j \neq i$ at the beginning of period t. In period τ_{min} j will expect to earn at most $.5\Pi_j^t$. However, Π_j^t is a function of the berry bushes that were not cleared during period $\tau_{min} - 1$. If ϕ is low, then the marginal value that j could have earned by clearing in $\tau_{min} - 1$ rather than employing some other strategy will be higher than $.5\Pi_j^t$, meaning that clearing in period $\tau_{min} - 1$ would be a dominant strategy. In fact, ϕ would need to be very high for this to not be the case - the value of all the berry bushes after all the harvesting decisions have been made would have to be less than half of the value of the berry bushes at the beginning of the next period. For $\phi = .0025$ this condition does not come close to attaining, indicating that a non-clearing subject j would be better off by responding to subject i's minimal clearing time by clearing even sooner than i. The rest of the argument proceeds via backwards induction.

Appendix B: Berry Island Subject Instructions

The following instructions are taken from the costly specialization treatment condition.

Introduction

Participation in this experiment is divided into several parts:

1. You will read these instructions.
2. You will complete a guided tutorial explaining the use of Second Life, a computer program used in the experiment
3. We will start the experiment, which involves using Second Life. The actual experiment will last for two hours.
4. At the end, you will be paid $7 for showing up, plus possibly an additional amount, depending on your earnings in the experiment. During this experiment, you will earn experimental dollars. For every $250 of experimental dollars you earn, you will receive $1 of real money.

If you have any questions, raise your hand.

Please do not write on the instructions. Also, at this time please put away any mobile devices which you may have.

Second Life

Second Life is a 3D simulation. You control a character called your avatar. Each avatar has a different name and will start with a gray shirt.. In the picture below, you see Peter Twieg (the avatar of one of the experimenters) with a gray shirt.

As a general important rule, do not use any feature of the interface that we do not describe in the instructions. If you do not follow this rule, we may have to disqualify you from participation and send you home.

2

Experimental Field

In the picture below, you can see the experimental field. You will be free to explore it as you please once the session begins, although your avatar training will take place in a different region.

Field Features:

The most important features of the field are the berry bushes and the campground. There are several dozen berry bushes scattered throughout the experimental field, and one campground in its center. The berry bushes are depicted in the screenshot below on the left, while the campground is depicted in the screenshot below on the right.

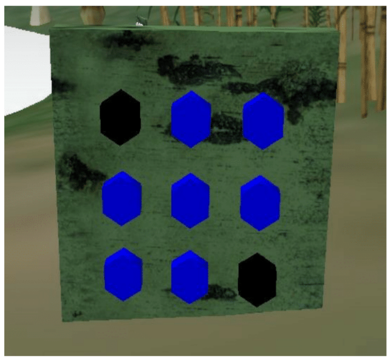

Berry bushes come in five varieties: Yellow, green, red, blue, and silver. The importance of these objects will be explained later in the instructions.

Using Second Life

1. To move around, you use the arrow keys. Go ahead and try this now. Additionally, you can hit the Page-Up key on your keyboard in order to jump.

2. In the bottom of the screen you see a blank text field that you can click on and type a message into. You can use this text field to chat with anyone who is within 20 meters of you, which is roughly 10 times the height of your avatar. To chat, type a message into this field and hit “Enter.” Go ahead and try this now – put a simple “hello” into chat. You should be able to see the messages of nearby avatars when this is done.

When the chat field is activated, you have to click somewhere else in the window before you can use the arrow keys to walk around again.

Berry Bushes

During the session, you can acquire seeds of five different colors. Seeds are collected from berry bushes of the corresponding color. In order to collect seeds from a berry bush, you must get close to it (within 5 meters) and click on a colored berry. When you do so, you will be given 4 seeds of that color and that berry will turn black. Black berries are considered to be depleted. Each seed you carry will be worth $1 to your total earnings in a round.

Each bush will have a maximum of nine colored berries at any given time. Note that the front and the back of a berry bush are the same: There are nine berries total, not nine on the front and nine on the back. This means the entire bush is depletable: If all the colored berries are picked, you will not be able to get any more seeds from that berry bush. At the beginning of each round, all berry bushes will start with between 1 and 3 seeds.

Depleted berries will slowly come back (regenerate) if the other berries on a given bush are not depleted. Every 5 seconds, the probability of a given depleted berry regenerating is:

P(r) = .02 * (number of undepleted berries) / 8

This rate of regenation is random and depends on the number of remaining berries. The following table illustrates the relationship between the number of berries remaining and the average rate at which they will regenerate.

# of undepleted berries	Expected time to regenerate one berry	Expected number of regenerations in 15 mins
0	Infinite	0
1	250	3.6
2	142	6.3
3	111	8.1
4	100	9
5	100	9
6	111	8.1
7	142	6.3
8	250	3.6
9	N/A	0

There are three important things to observe here:

(1) If a bush becomes completely depleted, it will no longer regenerate at all and will be useless for the rest of the given round.

(2) A bush will regenerate most-quickly when it has exactly four or five berries remaining on it.

(3) If harvested consistently when at its maximum regrowth rate, a berry tree will produce several times as many berries in a round than the amount it starts with.

To help keep track of this, all berry bushes will display text above them showing how long the expected regeneration rate is given the current number of berries they have.

Campground

In addition to generating money from seeds, you can acquire money from the campfire in the middle of the experimental field. As long as you're on the circular marble tile surrounding the campsite, you will continuously earn money at a rate of $12/minute. This may or may not be less lucrative than gathering seeds. This money will be added to the total you've earned at the end of the round and then reset for the next round.

HUD

In the top left corner of your screen, you see a translucent white panel. That is the HUD.

The HUD displays several pieces of information, from top to bottom:

1. It shows your name, how much time is remaining in the given part of the round, and how much you've earned in all previous rounds so far.
2. It shows how many of each type of seed you have. The numbers above each icon indicate the amount of that seed you currently own.

3. It shows how much money you're earning so far from various sources this round, broken down by source of earnings.

In addition, the HUD allows you to perform actions. By clicking on the icons corresponding to each seed color, a dialog menu will appear on your screen that will allow you to perform various actions on those seeds. During the session, you can do the following with seeds through these dialog menus.

1) Offering exchanges. You may offer a nearby player some amount of a type of seed that you have. In order to do this, pick the “Offer” button from the dialog menu, and then choose the name of the intended recipient from the list of nearby avatars that pops up. Next, a text box will pop up prompting you to enter an amount to offer. Type a number in that box and hit “Send” and your offer will be given to the recipient. If that recipient makes an offer to you within a 30-second time interval, then the offer will become a proposal: You will both receive a dialogue window giving the details of the proposal. If both parties accept the proposal within 30 seconds, then you will give your offer to the other person and that personal will give their offer to you.

Note that while a proposal is being considered, you cannot receive a proposal from anyone else. If you try to make a proposal with someone who is already considering one, then you will be told that the person is busy and the proposal will fail. If you choose the “Reject” option on the proposal's dialog window, then the proposal will fail and you will be able to entertain new ones.

2) Specializing in a certain type of seed. By choosing the “specialize” option, you will receive 10 seeds of the chosen color rather than 4 when a berry is picked. You can only have one specialization chosen at a time (specializing in a new color will undo the old specialization), but you may respecialize freely throughout the session. Note that you can only specialize in Red, Green, Yellow, or Blue berries – not Silver.

Next, you will be given a limited number of punishment tokens at the beginning of each round that you may or may not use by clicking on the black circle on your HUD. If you do so, you will be prompted to select from a list of names of nearby avatars, and then a quantity of punishment tokens to use – much like how one transfers seeds. Each punishment token you use will deduct $1 from the amount that you will earn for the round and will deduct $2 from the amount that the targeted subject will earn for the round. (Note: This may cause you to earn a negative amount for the given round!) You may use up to a maximum of 200 punishment tokens in any round, and the number above the punishment token will keep track of how many tokens you have remaining.

Experimental Procedure

Once started, the experiment will be divided into 8 rounds lasting 15 minutes each. During the rounds, you are free to explore the experimental field, acquire seeds, communicate, or do anything else explained in these instructions.

After the 4th round, there will be a brief break where some avatars will randomly be placed in new groups of 4. Also, the layout of the field will change at this time. The experiment will then restart for another 4 rounds.

At the end of each round, you will receive 1 experimental dollar for each seed inventory, plus 1 dollar for every 5 seconds spent at the campground, minus 1 dollar for every punishment token used *by* you, minus 2 dollars for every punishment token used *on* you. After this tallying is done, your inventory will be reset (all seeds will be taken away) and you'll be given a new set of punishment tokens. Additionally, your current specialization will be reset.

Also, all berry bushes will be reset for the start of the new round. They will all start with a new batch of 1-3 berries regardless of how many berries they had at the end of the previous round.

At the end of the entire experiment, you'll receive money based on the sum of how much you earned during each of the individual rounds of the session. Note that you may receive a negative amount in some rounds; however, you cannot earn a negative total across the entire experiment – the minimal score will be $0.

Very Important Rules about Second Life

As mentioned previously, the Second Life client makes it possible to do things you should not be doing in the experiment. As stated at the beginning, do not use any feature of the interface that we do not describe in the instructions.

Very Important Rules about Language

While using chat, please observe the following rules.

1. No profanity.
2. No communication in languages other than English.
3. No personal information (name, sex, age) that would allow other participants in the experiment to determine which person is controlling which avatar. This includes mentioning your computer number!
4. No spamming.

Very Important Rules about Laboratory Procedure

1. Do not attempt to look at anyone else's computer screen during this session.
2. Make sure any mobile devices are turned off and put away during this session.
3. If you have a question or if something does not appear to be working properly, please raise your hand and contact an experimenter.

Appendix C: Trust Networks Subject Instructions

These instructions are used in all of the treatment conditions.

Introduction

Participation in this experiment is divided into several parts:

1. You will read these instructions.
2. You will complete a guided tutorial explaining the use of Second Life, a computer program used in the experiment
3. We will start the experiment, which involves using Second Life. The actual experiment will last for two hours.
4. At the end, you will be paid $7 for showing up, plus possibly an additional amount, depending on your earnings in the experiment. During this experiment, you will earn experimental dollars. For every $500 of experimental dollars you earn, you will receive $1 of real money.

If you have any questions, raise your hand.

Please do not write on the instructions.

Second Life

Second Life is a 3D simulation. You control a character called your avatar. Each avatar has a different name and will have one of three different shirt colors. In the picture below, you see S13 Csn with a grey shirt.

As a general important rule, do not use any feature of the interface that we do not describe in the instructions. If you do not follow this rule, we may have to disqualify you from participation and send you home.

Experimental Field

In the picture below, you can see the experimental field. It is divided into three parts, with there also being a center area that is initially inaccessible. For the experiment, you will be placed in one of three groups, each of which will inhabit one part of the field (hereafter referred to as a "village.") Initially you will be unable to leave your own village.

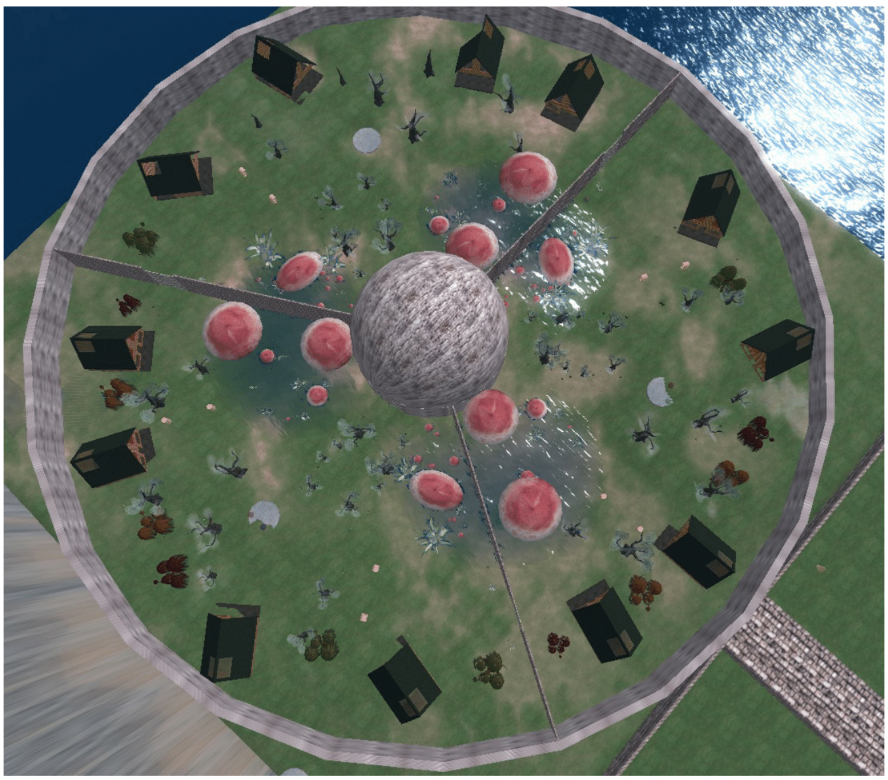

Village Features:

The most important features of the villages are the barrels and teleporters. Each village will have four barrels and a teleporter pad with two teleporters. These objects are depicted in the screenshots below.

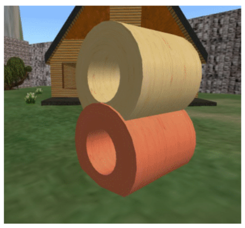

There are two types of barrels in your village: Ring barrels and circle barrels. Ring barrels will produce rings, and circle barrels produce circles. These items will explained in more depth later. The teleporter pads can be used to quickly travel between different villages during the session. You can also travel between villages by going through the middle area or the doors in the side walls once they become opened during the session.

Using Second Life

1. To move around, you use the arrow keys. Go ahead and try this now – run from wherever you currently are to one of the barrels, and then run to the teleporter pad in the middle of your village.

2. In the bottom of the screen you see an area that says "Click here to chat." You can use this text field to chat with anyone who is within 10 meters of you, which is roughly 5 times the height of your avatar. To chat, type a message into this field and hit "Enter." Go ahead and try this now – put a simple "hello" into chat. You should be able to see the messages of nearby avatars when this is done.

When the chat field is activated, you have to click somewhere else in the window before you can use the arrow keys to walk around again.

Acquiring Resources

During the session, you can acquire 4 different types of resources: circles, rings, triangles and diamonds. Circles, rings, and triangles all come in red, green, and blue, while diamonds are always black.

The colored barrels in the experimental field allow you to acquire new circles or rings of the same color as the base of the barrel. To do so, you must be within 5 meters of the barrel. Once you approach the barrel, it will say at the top which type of item it produces, and if you left-click on it using the cursor it becomes yours for as long as you stay within 5 meters of it. Only one person can own any one barrel at a time. When it becomes owned, the barrel will switch from being a pale color to a deeper one. Additionally, while it is owned, every second the barrel has a random chance of producing a circle or ring. When it does, the top of the barrel becomes colored as well, and at this point you can click on the barrel in order to acquire an object of the given type. Once the item has been acquired, the barrel will begin to produce a new item of that type.

Additionally, by producing an item of a given type, you will gain experience in collecting more of that item in the future. If you spend 240 seconds during a round producing items of a given type, you will be able to collect two of that item at once instead of one. If you spend 540 seconds producing items of a given type, you will be able to collect four of that item at once instead of one.

At this time, go ahead and find a barrel and we will practice using them in order to collect some resources. You must get close on them and click them, and then stay close and have it produce items.

Using Resources

Once you've acquired resources, you can perform various actions with them. Firstly, you can choose to transfer resources in your inventory to another player who is standing no more than 10 meters from you. You can also consume these resources, earning $10 for a ring or circle, $40 for a triangle, and $160 for a diamond.

You also have the option of combining certain resources. A circle and a ring of the same color can combine to form a triangle of that color. Two triangles of different colors can combine to form a black diamond. Diamonds cannot be combined and can only be transferred or consumed. Triangles and diamonds can only be acquired through combinations.

To perform any of these actions, you use your HUD, which is described on the next page.

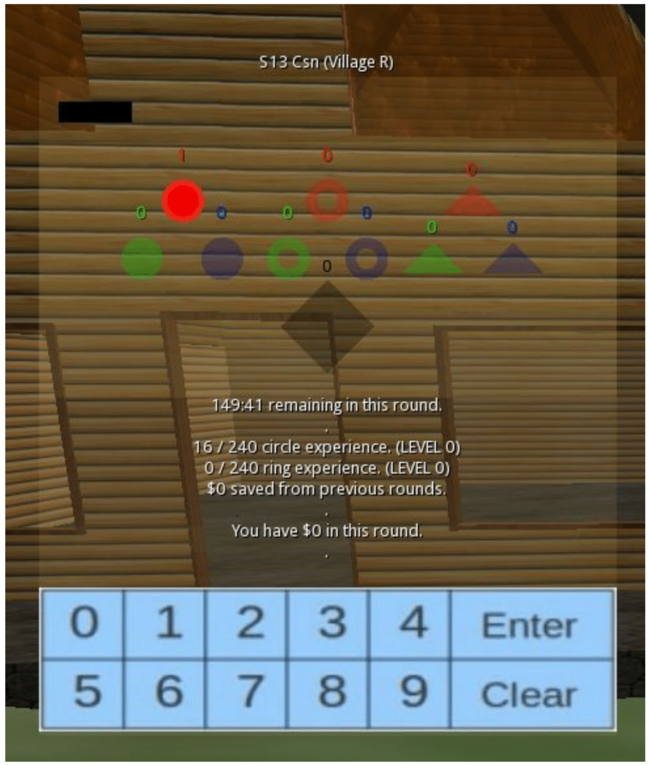

HUD

In the top left corner of your screen, you see a translucent white panel. That is the HUD.

The HUD displays several pieces of information.

1. It shows how many of each resource you have. The numbers above each icon indicate the amount of that resource you currently own.
2. It shows how much money you have.
3. It shows the amount of time remaining in the given round.
4. It displays your name and which village your avatar belongs to.

In addition, the HUD allows you to perform actions on resources. When you click on a circle, ring, or triangle, a menu pops up with options such as “consume”, ”combine”, “transfer”, and “improve.” The first three of these options can be performed on 1, 5, or 10 of the given resource at once using these options.

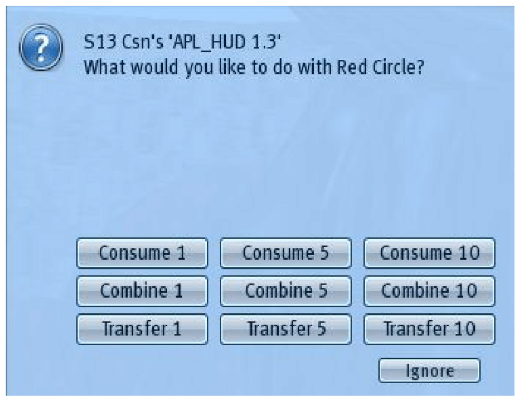

To combine two resources, click one of the resources, then click “combine.” If the resource is a ring or circle and you have a resource with which it can combine, it will do so automatically. Combining a triangle produces a menu asking which color triangle you wish to combine it with. Click on one of the options, but make sure you have a triangle of that color or the combination will fail.

To transfer a resource, click on the resource and then click “transfer.” You will see a menu. with a list of the names of nearby players. Click the name of the player to whom you want to transfer the resource. Players must be located within 10 meters of you to appear on this menu.

Lastly, you can also use the keypad at the bottom of the HUD to transfer money to another player. To do so, click on the number pad to enter the quantity you wish to transfer, and then choose “Enter.” A window with the names of nearby players will appear. Choose the name of the person to whom you want to transfer money. Once again, the player must be within 10 meters of you in order for this to work properly.

At this time, please try using the resources you've collected from the barrels in some way – try transferring them, or consuming them, or whatever makes you feel comfortable with the experimental interface.

Experimental Procedure

Once started, the experiment will be divided into 6 rounds lasting 15 minutes each. During the rounds, you are free to acquire resources, consume them, transfer them, etc.

At the end of each round, the amount of money you have will be added to a running total of what you've earned during the experiment, and then reset to 0 for the next round. Additionally, any leftover resources in your inventory will be automatically converted to cash. You cannot carry resources or money from one round to the next, but you will be paid at the end of the experiment based on how much your earn during each of the six rounds.

Additionally, your improvements will be reset at the end of each round and you will have to reacquire them in order to get the relevant bonuses.

Initially during the experiment, you will only be able to interact with members of your own village and will be unable to leave your village. After the second round, however, the doors between the villages will open, and also the teleporters in the middle of the villages will activate. You can use the teleporters once they've activated by walking up close to them (less than 2 meters) and clicking on them – you will be automatically transported to the teleporter pad in the village corresponding to the color of the teleporter you clicked on.

Keep in mind that although this allows you to interact with members of other villages, you will not be able to use barrels in these other villages. You can only use the barrels in the original village you inhabit.

Very Important Rules about Second Life

As mentioned previously, the Second Life client makes it possible to do things you should not be doing in the experiment. As stated at the beginning, do not use any feature of the interface that we do not describe in the instructions.

In particular, here are some of the things that you should not doo.

1. Travel in any other way than walking around.
2. Communicating in any other way than the chat field at the bottom of your screen.
3. Alter the appearance of your avatar.

We cannot make an exhaustive list, but please realize that the purpose of this experiment is not for you to discover obscure features of Second Life.

Very Important Rules about Language

While using chat, please observe the following rules.

1. No profanity.
2. No communication in languages other than English.
3. No personal information (name, sex, age) that would allow other participants in the experiment to determine which person is controlling which avatar. This includes mentioning your computer number!
4. No spamming.

Very Important Rules about Laboratory Procedure

1. Do not attempt to look at anyone else's computer screen during this session.
2. Make sure any mobile devices are turned off and put away during this session.
3. If you have a question or if something does not appear to be working properly, please raise your hand and contact an experimenter.

Bibliography

Acheson, James M, and Roy J Gardner. 2004. "Strategies, conflict, and the emergence of territoriality: the case of the Maine lobster industry." *American Anthropologist*, 296–307.

Aebischer, Nicholas J, Peter A Robertson, and Robert E Kenward. 1993. "Compositional analysis of habitat use from animal radio-tracking data." *Ecology*, 74(5): 1313–1325.

Al-Ubaydli, Omar, Daniel Houser, John Nye, Maria Pia Paganelli, and Xiaofei Sophia Pan. 2013. "The Causal Effect of Market Priming on Trust: An Experimental Investigation Using Randomized Control." *PloS one*, 8(3).

Al-Ubaydli, Omar, Kevin McCabe, and Peter Twieg. 2014. "Can more be less? An experimental test of the resource curse." *Journal of Experimental Political Science*, 1(01): 39–58.

Atlas, Stephen, and Louis Putterman. 2011. "Trust among the avatars: A virtual world experiment, with and without textual and visual cues." *Southern Economic Journal*, 78(1): 63–86.

Avey, James B, Bruce J Avolio, Craig D Crossley, and Fred Luthans. 2009. "Psychological ownership: Theoretical extensions, measurement and relation to work outcomes." *Journal of Organizational Behavior*, 30(2): 173–191.

Banerjee, Abhijit V, and Esther Duflo. 2009. "The experimental approach to development economics." *Annu. Rev. Econ.*, 1(1): 151–178.

Bardsley, Nicholas. 2008. "Dictator game giving: altruism or artefact?" *Experimental Economics*, 11(2): 122–133.

Bellman, Richard. 1957. *Dynamic Programming.* . 1 ed., Princeton, NJ, USA:Princeton University Press.

Berg, Joyce, John Dickhaut, and Kevin McCabe. 1995. "Trust, reciprocity, and social history." *Games and economic behavior*, 10(1): 122–142.

Blattman, Christopher, Julian Jamison, Tricia Koroknay-Palicz, Katherine Rodrigues, and Margaret Sheridan. 2016. "Measuring the measurement error: A method to qualitatively validate survey data." *Journal of Development Economics*, 120: 99–112.

Bochet, Olivier, Talbot Page, and Louis Putterman. 2006. "Communication and punishment in voluntary contribution experiments." *Journal of Economic Behavior & Organization*, 60(1): 11–26.

Brosig, Jeannette, Joachim Weimann, and Axel Ockenfels. 2003. "The effect of communication media on cooperation." *German Economic Review*, 4(2): 217–241.

Brown, David N, and Robert S Pomeroy. 1999. "Co-management of Caribbean Community (CARICOM) fisheries." *Marine Policy*, 23(6): 549–570.

Brown, Graham, Thomas B Lawrence, and Sandra L Robinson. 2005. "Territoriality in organizations." *Academy of Management Review*, 30(3): 577–594.

Camerer, Colin. 2011. "The promise and success of lab-field generalizability in experimental economics: A critical reply to Levitt and List."

Cárdenas, Juan-Camilo, and Elinor Ostrom. 2004. "What do people bring into the game? Experiments in the field about cooperation in the commons." *Agricultural systems*, 82(3): 307–326.

Carpenter, Jeffrey P. 2007. "The demand for punishment." *Journal of Economic Behavior & Organization*, 62(4): 522–542.

Charness, Gary, Ernan Haruvy, and Doron Sonsino. 2007. "Social distance and reciprocity: An Internet experiment." *Journal of Economic Behavior & Organization*, 63(1): 88–103.

Clay, Karen. 2006. "Squatters, Production, and Violence."

Cosmides, Leda, and John Tooby. 1992. "Cognitive adaptations for social exchange." *The adapted mind*, 163–228.

Crockett, Sean, Vernon L Smith, and Bart J Wilson. 2009. "Exchange and Specialisation as a Discovery Process*." *The Economic Journal*, 119(539): 1162–1188.

Demsetz, Harold. 1967. "Toward a theory of property rights." *The American economic review*, 57(2): 347–359.

Denant-Boemont, Laurent, David Masclet, and Charles N Noussair. 2007. "Punishment, counterpunishment and sanction enforcement in a social dilemma experiment." *Economic theory*, 33(1): 145–167.

DeScioli, Peter, and Bart J Wilson. 2011. "The territorial foundations of human property." *Evolution and Human Behavior*, 32(5): 297–304.

Desmet, Pieter, David De Cremer, and Eric van Dijk. 2011. "In money we trust? The use of financial compensations to repair trust in the aftermath of distributive harm." *Organizational Behavior and Human Decision Processes*, 114(2): 75–86.

Fafchamps, Marcel, and Bart Minten. 1999. "Relationships and traders in Madagascar." *The Journal of Development Studies*, 35(6): 1–35.

Fafchamps, Marcel, et al. 1998. "Market emergence, trust and reputation."

Fiedler, Marina, Ernan Haruvy, and Sherry Xin Li. 2011. "Social distance in a virtual world experiment." *Games and Economic Behavior*, 72(2): 400–426.

Field, Erica. 2007. "Entitled to work: Urban property rights and labor supply in Peru." *The Quarterly Journal of Economics*, 122(4): 1561–1602.

Fiore, Stephen M, Glenn W Harrison, Charles E Hughes, and E Elisabet Rutström. 2009. "Virtual experiments and environmental policy." *Journal of Environmental Economics and Management*, 57(1): 65–86.

Fischbacher, Urs. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental economics*, 10(2): 171–178.

Futuyma, Douglas J, and Gabriel Moreno. 1988. "The evolution of ecological specialization." *Annual Review of Ecology and Systematics*, 207–233.

Gintis, Herbert. 2007. "The evolution of private property." *Journal of Economic Behavior & Organization*, 64(1): 1–16.

Greif, Avner. 1993. "Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition." *The American economic review*, 525–548.

Greif, Avner. 2006. "History lessons: the birth of impersonal exchange: the community responsibility system and impartial justice." *The Journal of Economic Perspectives*, 20(2): 221–236.

Greiner, Ben, Mary Caravella, and Alvin E Roth. 2014. "Is avatar-to-avatar communication as effective as face-to-face communication? An Ultimatum Game experiment in First and Second Life." *Journal of Economic Behavior & Organization*, 108: 374–382.

Gürerk, Özgür, Andrea Bönsch, Lucas Braun, Christian Grund, Christine Harbring, Thomas Kittsteiner, and Andreas Staffeldt. 2014. "Experimental economics in virtual reality."

Harrison, Glenn W. 1989. "Theory and misbehavior of first-price auctions." *The American Economic Review*, 749–762.

Harrison, Glenn W, and John A List. 2004. "Field experiments." *Journal of Economic literature*, 42(4): 1009–1055.

Harrison, Glenn W, Ernan Haruvy, and E Elisabet Rutström. 2011. "Remarks on virtual world and virtual reality experiments." *Southern Economic Journal*, 78(1): 87–94.

Haushofer, Johannes, and Ernst Fehr. 2014. "On the psychology of poverty." *Science*, 344(6186): 862–867.

Heaton, Jeff. 2007. *Introduction to linden scripting language for Second Life.* Heaton Research, Inc.

Hoffman, Elizabeth, Kevin McCabe, and Vernon L Smith. 1996. "Social distance and other-regarding behavior in dictator games." *The American Economic Review*, 86(3): 653–660.

Innocenti, Alessandro. 2015. "Virtual reality experiments in economics."

Isaac, R Mark, and James M Walker. 1988. "Communication and free-riding behavior: The voluntary contribution mechanism." *Economic Inquiry*, 26(4): 585–608.

Janssen, Marco A, and Elinor Ostrom. 2008. "TURFS in the lab: institutional innovation in real-time dynamic spatial commons." *Rationality and Society*, 20(4): 371–397.

Janssen, Marco A, Robert Holahan, Allen Lee, and Elinor Ostrom. 2010. "Lab experiments for the study of social-ecological systems." *Science*, 328(5978): 613–617.

Karlan, Dean, Markus Mobius, Tanya Rosenblat, and Adam Szeidl. 2009. "Trust and social collateral." *The Quarterly Journal of Economics*, 124(3): 1307–1361.

Kimbrough, Erik O, and Bart J Wilson. 2011. "Geography and social networks in nascent distal exchange." *Journal of Institutional and Theoretical Economics JITE*, 167(3): 409–433.

Kimbrough, Erik O, Vernon L Smith, and Bart J Wilson. 2008. "Historical Property Rights, Sociality, and the Emergence of Impersonal Exchange in Long-Distance Trade." *American Economic Review*, 98(3): 1009–1039.

Kimbrough, Erik O, Vernon L Smith, and Bart J Wilson. 2010. "Exchange, theft, and the social formation of property." *Journal of Economic Behavior & Organization*, 74(3): 206–229.

Kühberger, Anton. 1998. "The influence of framing on risky decisions: A meta-analysis." *Organizational behavior and human decision processes*, 75(1): 23–55.

Leacock, Eleanor. 1954. "The Montagnes 'Hunting Territory' and the Fur Trade." *American Anthropologist*, 56(5): 1–59.

Leeson, Peter T. 2007. "Anarchy, monopoly, and predation." *Journal of Institutional and Theoretical Economics JITE*, 163(3): 467–482.

Levitt, Steven D, and John A List. 2007. "What do laboratory experiments measuring social preferences reveal about the real world?" *The journal of economic perspectives*, 21(2): 153–174.

Lévy-Garboua, Louis, Hela Maafi, David Masclet, and Antoine Terracol. 2012. "Risk aversion and framing effects." *Experimental Economics*, 15(1): 128–144.

Liberman, Varda, Steven M Samuels, and Lee Ross. 2004. "The name of the game: Predictive power of reputations versus situational labels in determining prisoners dilemma game moves." *Personality and social psychology bulletin*, 30(9): 1175–1185.

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty impedes cognitive function." *science*, 341(6149): 976–980.

Mohr, Carl O. 1947. "Table of equivalent populations of North American small mammals." *American midland naturalist*, 223–249.

Moore, Dana, Michael Thome, and Karen Haigh. 2009. *Scripting your world: the official guide to second life scripting.* John Wiley & Sons.

Nicklisch, Andreas, and Tobias Salz. 2008. "Reciprocity and status in a virtual field experiment."

Nikiforakis, Nikos, and Hans-Theo Normann. 2008. "A comparative statics analysis of punishment in public-good experiments." *Experimental Economics*, 11(4): 358–369.

North, Douglass C. 1990. *Institutions, institutional change and economic performance.* Cambridge university press.

North, Douglass . 1991. "Institutions STÖR." *The Journal of Economic Perspectives*, 5(1): 97–112.

Olson, Mancur. 1993. "Dictatorship, Democracy, and Development." *American Political Science Review*, 87(03): 567–576.

Oosterbeek, Hessel, Randolph Sloof, and Gijs Van De Kuilen. 2004. "Cultural differences in ultimatum game experiments: Evidence from a meta-analysis." *Experimental Economics*, 7(2): 171–188.

Ostrom, Elinor. 1990. *Governing the commons: The evolution of institutions for collective action.* Cambridge university press.

Ostrom, Elinor. 2000. "Collective action and the evolution of social norms." *The Journal of Economic Perspectives*, 14(3): 137–158.

Ostrom, Elinor, and James Walker. 2003. *Trust and reciprocity: Interdisciplinary lessons for experimental research.* Russell Sage Foundation.

Prediger, Sebastian, Björn Vollan, and Markus Frölich. 2011. "The impact of culture and ecology on cooperation in a common-pool resource experiment." *Ecological Economics*, 70(9): 1599–1608.

Pyke, Graham H. 1984. "Optimal foraging theory: a critical review." *Annual review of ecology and systematics*, 523–575.

Rietz, Thomas A, Roman M Sheremeta, Timothy W Shields, and Vernon L Smith. 2013. "Transparency, efficiency and the distribution of economic welfare in pass-through investment trust games." *Journal of Economic Behavior & Organization*, 94: 257–267.

Slater, Mel, Martin Usoh, and Anthony Steed. 1994. "Depth of presence in virtual environments." *Presence: Teleoperators & Virtual Environments*, 3(2): 130–144.

Smith, Adam. 1776. *The wealth of nations [1776].* na.

Smith, Vernon L. 1982. "Microeconomic systems as an experimental science." *The American Economic Review*, 923–955.

Smith, Vernon L. 2003. "Constructivist and ecological rationality in economics." *The American Economic Review*, 93(3): 465–508.

Spann, Martin, Oliver Hinz, Il-Horn Hann, and Bernd Skiera. n.d.. "Decision making in virtual worlds: An experimental test of altruism, fairness and presence." *Decision Making*, 1: 1–2010.

Thom-Santelli, Jennifer, Dan R Cosley, and Geri Gay. 2009. "What's mine is mine: territoriality in collaborative authoring." 1481–1484, ACM.

Tversky, Amos, and Daniel Kahneman. 1985. "The framing of decisions and the psychology of choice." In *Environmental Impact Assessment, Technology Assessment, and Risk Analysis*. 107–129. Springer.

Wason, Peter C. 1968. "Reasoning about a rule." *The Quarterly journal of experimental psychology*, 20(3): 273–281.

Wilde, Louis L. 1981. "On the use of laboratory experiments in economics." In *Philosophy in Economics*. 137–148. Springer.

Wilson, Bart J, Taylor Jaworski, Karl E Schurter, and Andrew Smyth. 2012. "The ecological and civil mainsprings of property: an experimental economic history of whalers' rules of capture." *Journal of Law, Economics, and Organization*, 28(4): 617–656.

Zizzo, Daniel John. 2010. "Experimenter demand effects in economic experiments." *Experimental Economics*, 13(1): 75–98.

Biography

Peter Twieg graduated from Theodore Roosevelt High School, Kent, Ohio, in 2002. He received his Bachelor of Arts from Reed College in 2006 and his Master of Arts from George Mason University in 2009. As of this writing, he is currently employed as a data scientist at Intuit.