

UNIVERSITY OF CALIFORNIA
Santa Barbara

Three Essays in Behavioral Economics

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Economics

by

Daniel Kaiser Saunders

Committee in Charge:

Professor Zachary J. Grossman, Chair

Professor Gary B. Charness

Professor Theodore C. Bergstrom

December 2014

The Dissertation of
Daniel Kaiser Saunders is approved:

Professor Gary B. Charness

Professor Theodore C. Bergstrom

Professor Zachary J. Grossman, Committee Chairperson

September 2014

Three Essays in Behavioral Economics

Copyright © 2014

by

Daniel Kaiser Saunders

To Kelly, for always supporting me,
and to my father, for showing me the way.

Acknowledgements

First, I would like to thank my advisor, Zack Grossman, for his encouragement, guidance and mentoring over the years. I would also like to thank my committee members, Gary Charness and Ted Bergstrom, for reading the many drafts of my papers and providing insightful comments and suggestions. I am grateful to my colleague, John Hartman, for sharing his experimental data for this project.

I would especially like to thank Ted Turocy, who generously shared his time and his Python code for tracing the Quantal Reponse Equilibria. Truly, this would not have been possible without his many years of yeoman's work on this topic. Last but not least, I would like to thank my gracious co-author, Ayana E. Johnson, for the opportunity to work on a collaborative publication.

Curriculum Vitæ

Daniel Kaiser Saunders

Contact

Email: saunders@econ.ucsb.edu

Website: <http://econdan.com>

Education

2014 (Expected) PhD in Economics, University of California, Santa Barbara.

2009 Master of Arts in Economics, University of California, Santa Barbara.

2007 Bachelor of Arts in Economics, Mathematics, and Spanish, University of Kansas.

Research Interests

Behavioral and experimental economics, environmental and natural resource economics

Experience

- 2013 – 2014 Lecturer at the Martin V. Smith School of Business, California State University, Channel Islands
- 2007 – 2013 Teaching assistant for undergraduate and graduate Courses, University of California, Santa Barbara
- 2005 – 2007 Instructor for college algebra in the Kansas Algebra Program, University of Kansas
- 2010 Research Assistant for Zachary Grossman, University of California, Santa Barbara

Working Papers

D.K. Saunders *Errors and altruism in threshold public good.*
2014

D.K. Saunders *The trouble with trembles: Reverberating noise in the market entry game.* 2014

Publications

A.E. Johnson and D.K. Saunders *Time preferences and the management of coral reef fishers..* Ecological Economics; April 2014, Vol. 100, pp. 130-139.

Awards and Honors

Outstanding Teaching Assistant Award. UCSB 2012

Nominee for GSA Excellence in Teaching Award. UCSB
2013

Conference Participation

May 2014 Bay Area Behavioral/Experimental Economics Workshop, Stan-
ford CA

October 2013 North-American ESA Conference, Santa Cruz CA

November 2012 North-American ESA Conference, Tucson AZ

Abstract

Three Essays in Behavioral Economics

Daniel Kaiser Saunders

This dissertation concerns two key topics in behavioral economics: bounded rationality and hyperbolic time preferences. While these two topics do not have much in common, per se, the analyses presented here do contain unifying themes. Each chapter focuses on issues pertaining to public economics. The environments are often characterized by strategic interaction, in that each individual's choices affects the welfare of others. In such cases, economic efficiency is by no means guaranteed.

The first two chapters examine environments with multiple Nash equilibria, which limits predictions. I utilize Quantal Response Equilibrium, which relaxes the assumption of pure best-response in favor of noisy best response. This allows me to make comparative static predictions across group-size, information completeness, payoff-type, and state of the world; predictions that are problematic under pure best-response. I conduct laboratory experiments for the threshold game (chapter one) and the market entry game (chapter two). I find that much of the seemingly anomalous behavior observed in the data is easily explained using QRE. This suggests that quantal response is a useful alternative assumption to

best-response in coordination problems, and, in particular, logit QRE is a valuable equilibrium model of coordination failure with simple economic intuition.

The third chapter advances a simple $\beta - \delta$ model of quasi-hyperbolic time preferences. In this chapter, the theoretical model is used to construct the key variables from our data. This data was collected in the field by my co-author, when she visited the islands of Bonaire and Curaçao in 2010. While much of the data is survey response or demographic related, the last part of each interview includes a paid experiment to elicit time preferences over various time horizons. We find that, controlling for key demographic factors, decreased patience (lower δ) and present bias ($\beta < 1$) are associated with an lower willingness to adopt marine reserve restrictions, but they have no effect on attitudes about gear restrictions.

Professor Zachary J. Grossman
Dissertation Committee Chair

Contents

Acknowledgments	v
Curriculum Vitæ	vi
Abstract	ix
List of Figures	xiv
List of Tables	xv
Introduction	1
1 Explaining pledges to threshold public goods with logit equilibrium	6
1.1 Introduction	6
1.2 Equilibria in the Threshold Game	12
1.2.1 Nash and Bayes-Nash Equilibria	15
1.2.2 Quantal Response Equilibria	21
1.2.3 Quantal Response Equilibrium Predictions	27
1.3 The Experiment	38
1.3.1 Analysis Across Individuals	40
1.3.2 Analysis Across Groups	45
1.3.3 Quantal Response Equilibrium Results	49
1.4 Discussion	57
1.5 Conclusion	68
2 Noisy best response and scale illusion in the market entry game	72
2.1 Introduction	72

2.2	Equilibria in the Market Entry Game	78
2.2.1	Nash and Bayes-Nash Equilibria	78
2.2.2	Quantal Response Equilibria	82
2.2.3	Quantal Response Equilibrium Predictions	86
2.3	The Experiments	89
2.3.1	Experiment I	91
2.3.2	Experiment II	96
2.4	Discussion	115
2.5	Conclusion	119
3	Time preferences and the management of coral reef fisheries	121
3.1	Introduction	121
3.2	Methods	124
3.2.1	Socioeconomic interviews	124
3.2.2	Eliciting Time Preferences	125
3.2.3	Calculating discount factors and present bias	127
3.2.4	Eliciting management preferences	130
3.2.5	Revealed management preferences	131
3.2.6	Data analysis	132
3.2.7	Risk-aversion	134
3.3	Results and Discussion	135
3.3.1	Interviewee demographics	135
3.3.2	Fisher and diver IDFs	138
3.3.3	Distribution of IDFs	141
3.3.4	Comparison with previous IDF studies	142
3.3.5	Present bias	143
3.3.6	Management scores	145
3.3.7	Principal component analysis of management scores	148
3.3.8	Restricting their own resource use	152
3.3.9	Key findings	153
3.3.10	Policy implications	154
	Conclusion	159
	Bibliography	161
	Appendices	169
A	Appendix to Chapter 1	170
A.1	Symmetric Bayes-Nash Equilibria	170

B	Appendix to Chapter 2	176
B.1	Theorems and Proofs	176
C	Appendix to Chapter 3	183
C.1	Time preference experiment	183
C.2	Survey questions about management of coral reef resources	186

List of Figures

1.1	QRE: Individual Contributions	28
1.2	QRE: Total Contributions	30
1.3	Group-Level QRE Predictions	32
1.4	Individual-Level QRE Predictions	35
1.5	Histograms: Individual Contributions	41
1.6	Histograms: Total Contributions	43
1.7	Mean/SD Individual Contributions	44
1.8	Total Contributions by Group across Rounds	47
1.9	Total Contributions by Group across Rounds	52
1.10	Mean Individual Payoffs by Contribution	62
2.1	Individual Entry Rates	95
2.2	QRE: Principal Branch	100
2.3	Average Entry Relative to Capacity	101
2.4	Low Entry Rates	107
2.5	High Entry Rates	109
2.6	Average Entry Rates	113
2.7	Individual Entry Rates	115
3.1	Distributions of individual discount factors	140
A.1	Generic Payoff Functions	171
A.2	Sets of symmetric Bayes-Nash equilibria	172

List of Tables

1.1	Bayes-Nash Equilibria	20
1.2	Serial Correlation in Total Contributions	48
1.3	Maximum Likelihood Estimates (all 20 rounds)	49
1.4	Estimated and Predicted Levels for Groups	51
1.5	Information and Group-size Effects for Groups	53
1.6	Estimated and Predicted Levels for Individuals	55
1.7	Information, Group-size, and Payoff-Type Effects for Individuals	56
2.1	Experiment I: Average Entry by Treatment	93
2.2	Experiment II: Payoffs	97
2.3	Nash Equilibria	99
2.4	Observed Entry Rates Across States	102
2.5	Predicted Entry Rates Across States	102
2.6	Test of BNE restriction	104
2.7	Average Observed Benefit to Entry	105
2.8	Tests for Persistence	110
2.9	Logit QRE versus Logit Regression	118
3.1	Price list switch points with implied discount factors	128
3.2	Summary of demographic factors and management scores	137
3.3	Marginal effect of discount factors on management scores	146
3.4	Marginal effect of present bias on marine reserve scores	147
3.5	Marginal effect of discount factors on the principal components	150
3.6	Marginal effect of present bias on the third principal component	151
A.1	Comparing Expected Payoffs at Discontinuities	173
A.2	Comparing Expected Payoffs at Discontinuities	174

Introduction

Behavioral models serve a useful role as a complement to standard economic models. There are important contexts in which a behavioral model or insight provides the greatest predictive power. This dissertation explores three such contexts: the case of a threshold public good, the problem of market entry (a congestible public good), and the case managing an exhaustible natural resource. Public goods problems are just one example of situations involving strategic interaction, where individual best response behavior does not guarantee overall economic efficiency. This supports the case for policy interventions aimed at increasing welfare. Such policies require a correct diagnosis of the underlying failure, for which I turn to behavioral models.

Chapter 1 focuses on the threshold public goods game, in which individuals simultaneously choose a level of pledged contributions to a unitary public good. If total pledges exactly meet an exogenous threshold that is common knowledge, then all pledges are binding and the public good is produced. I conduct a two-

by-two experiment varying group-size and information regarding the state of the world. Small groups have four members, while large groups have eight. Players are endowed with one of two possible values to the public good, drawn independently each round with equal probability. Under complete information, subjects knew the entire state of the world, while under incomplete information each player has private knowledge of her own type, and she knows the binomial distribution from which others' types are drawn.

It is hard to make predictions across these treatments using Nash or Bayes-Nash equilibrium concepts since there are typically infinitely many equilibria in this game. I elect to use logit quantal response equilibrium, which provides a straightforward method for making predictions across these treatments. The model under-predicts average contributions and the frequency with which the threshold is met, called the provision rate; although it does a better job predicting changes in these variables across treatments.

One potential improvement to this research would be to incorporate a structural estimate of social preferences. Given the stochastic nature of QRE, a likelihood function may be constructed to estimate the degree of decision error and altruism, simultaneously. This would likely remove the biased under-predictions. This would also help place the paper with in the public goods literature, where other-regarding preferences play a crucial role. As is, logit equilibrium has con-

siderable statistical power; supporting notion of quantal response behavior in the face of coordination problems.

Chapter 2 examines the market entry game. The environment is defined as follows: players choose between a risk-less “outside” option, where payoffs are fixed, and a risky “inside” option, where payoffs are decreasing in the total number of players who choose in, and there is some level of total entry at which payoffs fall below the outside option. This closely resembles the classic market entry problem with a finite capacity for profitable firms. Notice that the market, itself, is a congestible good. Another way to conceive of this game is as a binary analogue to the tragedy of the commons, and one commonly studied application is the classic highway congestion problem. Simply define the outside option as a lengthy beltway with no traffic and the inside option as a short bridge with travel times are increasing in the total number of travelers.

This chapter offers new analysis of data from a previous experiment by my colleague John Hartman, as well as an experiment of my own. Combining our data sets allows us to control for several experimental design choices and interpretations. Treatments include manipulating payoffs through either a toll or through affine re-scaling. I also study the effects of group-size and information completeness. I find that the QRE is able to capture key features of the data, and I demonstrate the econometric value-added by this approach by showing that the

logit QRE is able to predict a difference-in-difference effect across payoff-type and group-size that a standard logit regression will miss. This is because of the dual-causality problem implied by quantal choice in a strategic setting. Specifically, I learn that a model in which the reverberation of noise across individuals is fully accounted requires fewer degrees of freedom and provides greater statistical power than a model that assumes each player ignores the noise of others.

Chapter 3 presents my work with Ayana E. Johnson. We study the conservation attitudes of fishers and divers on the islands of Bonaire and Curaçao. These two professions represent the bulk of industry on these islands, and each one relies heavily upon the local coral reef fisheries. A paid field experiment was conducted in order to elicit revealed discount rates in a three period framework. Payments in the present period were distributed with a front-end delay to fully separate risk and time preferences. We motivate our analysis with a simple $\beta - \delta$ model, which amounts to measuring the changes in measured discount rates between the first two periods and the last two period. We are able to identify the individual discount rates and present bias in the data. Controlling for a host of demographic controls, we find that decreased patience (lower δ) or present bias ($\beta < 1$) reduces willingness to use marine reserves, but has no effect on preference for gear restrictions. This makes a great deal of sense since marine reserves rep-

resent inter-temporal conservation, while gear restrictions are design to mitigate a contemporaneous externality generated by fishers.

All of these works utilize theoretical models to motivate an empirical analysis using experimental data. Each chapter considers a coordination failure regarding a different public choice problem. In each example, standard economic models provide a useful basis for studying the problem. However, the addition of behavioral assumptions prove useful in increasing statistical power and predicting so-called anomalies in individual decision-making.

Chapter 1

Explaining pledges to threshold public goods with logit equilibrium

1.1 Introduction

Many fundraisers make use of a provision point: a predetermined threshold that total pledged contributions must exceed in order to produce a good. Notable examples include unitary public goods, such as a bridge or park, or low-marginal-cost information goods, such as those listed on www.kickstarter.com or www.indiegogo.com. These online implementations operate on a unique scale

compared to previous methods, both in the number of contributors and the value of contributions. More than 2,241,475 people pledged to at least one Kickstarter project in 2012, generating \$319,786,629 in total pledges. This led to 18,109 successfully funded projects. At least 17 separate projects raised in excess of one million dollars each. Yet, some 56% of all projects were not successfully funded.¹ To what extent the success rate of funding reflects the provision rate of economically efficient projects remains an open question.

As online threshold mechanisms continue to grow, it is increasingly important to understand the unique challenges that they present. There are two key features that distinguish these online technologies from other fundraising methods: (i) uncommonly large groups and (ii) little information regarding other contributors' incentives. Conventional thinking suggests that increased payoff uncertainty will reduce provision. Likewise, it seems intuitive that larger groups will have a more difficult time reaching the threshold.² Yet, it is not clear how these forces will interact in practice. Assuming total pledges are near the threshold, upward trembles in pledging will increase the likelihood of success, while downward trembles will reduce the odds. Therefore, individual noisiness can have a substantial

¹All estimates are taken from the Kickstarter website.

²This assumes the threshold is scaled in proportion to total benefits, so that the relative difficulty of the coordination problem is maintained.

effect on overall welfare, driven by the discontinuity of everyone's payoffs at the threshold.

Such strong strategic interaction means that the cost of even a small degree of coordination failure may be substantial. There are actually two coordination problems underlying this game. First, there is a standard coordination problem³ between the dominated equilibria with under-contribution and the efficient equilibria with total contributions equal to the threshold. The latter set of equilibria satisfy the undominated, trembling-hand-perfect equilibrium refinement (Bagnoli and Lipman, 1989). Second, there remains a conflicting-interest coordination problem among these efficient equilibria, similar to a game of chicken (Kagel and Roth, 1997).⁴ Many combinations of contributions are capable of achieving the threshold; each one implying a different distribution of net benefits to the players. The tension of the game lies between each player's incentive to pledge as little as possible, while still meeting the threshold as a group.

³In the pure coordination game, there are two equilibria, one of which is both Pareto efficient and Pareto improving. This is the simplest coordination problem, since both players have the same preferences over equilibrium selection.

⁴In the game of chicken, there are two equilibria, and each player prefers a different equilibrium. Interestingly, each equilibrium involves players taking opposite actions. Hence, players' decisions are strategic substitutes, and this is an anti-coordination game. In the threshold game, players' pledges are strategic complements below the threshold, but they are strategic substitutes at the threshold.

In their experiment, Bagnoli and McKee (1991) observe that larger groups converge to equilibrium slower than smaller groups under complete information.⁵ While this finding is consistent with the intuition that coordination difficulty increases as group-size grows, standard equilibrium notions are silent about this process. Likewise, Marks and Croson (1999) study the role of payoff uncertainty⁶ using provision points, but they find no significant effects on average contributions or the provision rate. By contrast, I use the set of symmetric, pure-strategy Bayes-Nash equilibria with private information to demonstrate how incomplete information might reduce provision; although, the problem of multiple equilibria prevents crisp predictions here too.

These empirical findings across group-size and information treatments may be consistent with the many equilibria under best-response, but the usefulness of that approach is limited by its inability to make precise predictions. I use logit quantal response equilibrium as an alternate explanation for these effects, and I motivate

⁵Bagnoli and McKee (1991) find slower convergence to equilibrium play for large groups. In addition, Goeree et al. (2005a) study the Volunteer's Dilemma, a parameterization of the provision point in which one player can unilaterally produce the public good. They derive and reject a mixed-strategy Nash equilibrium explanation of group size effects under complete information, and suggest noisy best response and inequity aversion as alternative explanations.

⁶The authors provide each player with private information regarding her valuation of the public good, but they provide no information, whatsoever, regarding others' valuations. Subjects were able to partially observe outcomes between experimental rounds, so it is possible that they could update upon prior beliefs regarding others' payoff types. However, the paper makes no attempt to model this process.

this approach with the dynamic process of noisy directional learning developed by Anderson et al. (2004).

While existing experiments document some of these treatment effects, they lack the unified design required for making clean comparisons across treatments. I conduct a new two-by-two experiment that varies group-size and payoff uncertainty, in order to examine these issues with greater clarity. Unlike Bagnoli and McKee (1991), the threshold is scaled in proportion to the number of players so as to preserve the difficulty of the coordination problem. It would be prohibitively expensive to conduct a paid experiment on the scale of online mechanisms, so I examine the more modest comparative static as group-size grows from four to eight players. For simplicity, payoff heterogeneity takes the simple form of a Bernoulli random variable, with high and low payoffs, implying a binomial distribution for the state of nature. Under the complete information treatment, each subject knows the full state of the world, while under the incomplete information treatment, she only knows her own payoff-type and the distribution from which others' types are drawn. This definition of payoff uncertainty contrasts sharply with Marks and Croson (1999), who offer subjects no information about others' payoff types, and the Bayesian game framework allows for theoretical tractability with both equilibrium concepts.

I find that increased group-size reduces the provision rate and expected total pledges (relative to the threshold) under complete information, but it has no effect under incomplete information. Interestingly, increased information increases average total pledges for small groups, but not for large groups, and increased group-size reduces the variance of total pledges, but only under complete information. At the individual level, I find that increased information increases the average contribution of high types in small groups, while it lowers the average contribution of low types in large groups. Lastly, increased group-size reduces the average contribution of high types under complete information. Logit QRE correctly predicts the sign of each of these results; comparative statics that were not readily available under pure best-response.

While the QRE model has substantial success accounting for comparative statics across treatments, it does not perform well when predicting levels within treatments. The model systematically under-predicts average contributions and the provision rate. Being a public goods game, it is natural to consider behavioral preferences for altruism or fairness, which would naturally increase the level of contributions. However, such models are of limited use with strict best-response, as they can rationalize many combinations of contributions at or above the threshold, including instances where some individuals contribute more than their values. Moreover, I find that total contributions exactly equaled the threshold only about

4% of the time, while they over-shot and under-shot the threshold at rate of 54% and 42% respectively. Exact hits to the efficient Nash equilibria are rare; though, there is a distinct bias above the threshold. The former is consistent with decision error while the latter with altruism. Thus, a model that incorporates both features may ultimately prove best.

Strong assumptions such as strict best response are useful in pursuit of simple intuition and mathematical tractability. Relaxing such assumptions is equally useful when it serves that same purpose, and I will show that the threshold game is one such instance. The remainder of the paper is structured as follows. Section 1.2 defines the threshold game and presents the predictions of various equilibrium concepts. Section 1.3 describes the experimental design, the data, and the estimation results for the model. Section 1.4 discusses the consequences of incorporating social preferences. Section 1.5 concludes with remarks for future research.

1.2 Equilibria in the Threshold Game

In this section I will present a description of the threshold game along with some background on research in this area. I will review the predictions of Nash under complete information, I will present some results for the symmetric, Bayes-Nash equilibria under incomplete information, and I will motivate and define the logit quantal response equilibria.

The threshold game represents a public choice regarding the production of a unitary public good with a known, exogenous minimum fundraising threshold. Each player has a fixed endowment of the private good with which to make pledges to the public good, and each player receives some private benefit if the public good is produced. Players simultaneously choose pledges. If total pledges fall short of the threshold, all pledges are refunded. If pledges equal the threshold, they are binding, and each player receives her benefit minus her pledge. If total contributions exceed the threshold, the public good is produced, and all excess contributions are refunded in proportion to pledges.

Early experiments on the provision point mechanism date back to Marwell and Ames (1980), who included a provision point treatment in their series of experiments documenting the absence of free-rider behavior in the laboratory. Subsequent research examined how different parameterizations of the threshold or individual benefits might affect behavior. Croson and Marks (2000) identified the step return (SR), the ratio of total benefit to the threshold (total cost), as analogous to the marginal per capita return in controlling for group-size effects. The MPCR was identified by Isaac and Walker (1988) to explain residual group-size effects in experiments with linear public goods. As it happens, the MPCR and the SR coincide at the threshold because the average marginal return per player

must equal the overall group return.⁷ Therefore, the threshold must be adjusted in proportion to average total benefit order to isolate pure group-size effects.

Bagnoli and Lipman (1989) established the equilibrium refinement necessary and sufficient for efficient outcomes in the provision point mechanism; both for a unitary and a multi-unit public good. Specifically, they observed that the undominated, trembling-hand-perfect Nash equilibrium refinement coincides with the subset of efficient equilibria at the threshold in the unitary case. In a follow-up study by Bagnoli and McKee (1991), the authors find support for this refinement, while Bagnoli et al. (1992) find strong evidence rejecting a similar refinement for the multiple-unit case. The authors also choose to use proportional refunds. This simplifies the notion of economic efficiency to the provision rate, i.e. the probability of meeting or exceeding the threshold, while avoiding any perverse incentives for over-pledging.⁸

Little is understood about the set of Bayes-Nash equilibria under incomplete information. This is due, in part, to the mathematical intractability of pure best-response in this highly discontinuous environment. In the next subsection, I

⁷Below the threshold, the marginal per capita return is zero, while above the threshold it is negative. At some infinitesimal $\varepsilon > 0$ short of the threshold, the marginal contribution will be pivotal, and the gross return to the group as a whole is fixed at $\sum_i v_i/T$. While individuals will have heterogeneous private returns, the average (per capita) return must equal the group return since total costs and benefits are fixed.

⁸This is not meant to imply that rebates have no effect on choices. Marks and Croson (1998) find that rebates affect the variance of contributions. As QRE makes predictions about the entire distribution of contributions, it may explain these refund effects.

will preview the parameterization of the game used for the experiment in order to numerically solve for the set of symmetric Bayes-Nash equilibria. For this parameterization, aggregate contributions must equal to the threshold in some state of nature, in equilibrium; leading to average provision rates between zero and one. However, this solution concept again yields infinitely many equilibria. By contrast, I demonstrate that logit equilibrium has the potential to account for noisy decision-making, generate predictions across all experimental treatments and variables, and explain any systemic patterns observed at the individual and group level.

1.2.1 Nash and Bayes-Nash Equilibria

Now I will review the set of Nash and Bayes-Nash equilibria. While these equilibria may vary across treatments, there are some unifying features. First, both concepts admit infinitely many equilibria, making it is difficult to predict outcomes. Additionally, the probability that contributions meet or exceed the threshold, called the provision rate, depends upon the equilibrium selected. There are two sets of equilibria under complete information: an efficient set where total contributions perfectly match the threshold with probability one, and a set of dominated equilibria where the probability of reaching the threshold is zero. The various provision rates predicted by the Nash equilibria in the Bayesian game with

private information lie in between these two extremes. There are multiple sets of equilibria, and each equals the threshold with probability one in some state of nature, yielding average provision rates between zero and one across states. This approach cannot make more specific predictions without assuming away the coordination problem that real players typically fail to resolve.

The structure of the model is as follows. There are n players, indexed as $i \in \{1, \dots, n\}$. Each player has a potential private benefit, v_i , for a unitary public good. Payoff functions are defined over pledged contributions $c_i \in [0, w_i]$ where w_i is player i 's endowment. If total pledges fall short of an exogenous and commonly known threshold, T , then the public good is not produced and all pledges are refunded. If the threshold is surpassed, then the public good is produced and total excess contributions are refunded in proportion to each player's share of total pledges. I express this in strategic notation below, with payoffs normalized to net gains for simplicity.

$$u_i(c_i, c_{-i}) = \begin{cases} v_i - T \left(\frac{c_i}{c_i + c_{-i}} \right) & \text{if } c_i + c_{-i} \geq T \\ 0 & \text{if } c_i + c_{-i} < T \end{cases} \quad \text{where } c_{-i} = \sum_{j \neq i} c_j \quad (1.1)$$

I assume $\sum_i v_i \geq T$ throughout, as the alternative case is of little interest. A strategy profile achieves the threshold, $\sum_i c_i = T$, if-and-only-if it is an undominated, trembling-hand-perfect Nash equilibrium (Bagnoli and Lipman, 1989). For

most parameterizations, a continuum of equilibria survive this refinement, and each equilibrium affords a different distribution of net benefits to players. There remains a conflicting-interest coordination problem among these efficient equilibria, which provides a continually destabilizing force against perfect coordination, as some players will always have an incentive to nudge the vector of contributions away from its current allocation. There also exists a set of dominated equilibria where the threshold is never achieved. In fact, total contributions are so low that no individual can deviate to a higher pledge and meet the threshold in an individually rational way.

Less is understood about the Bayes-Nash equilibria under private information, where each player, i , has private knowledge of v_i but only know the distribution of v_{-i} . The symmetric Bayes-Nash equilibria are a sensible starting point for this analysis since all players of a given payoff type face an identical decision problem. In addition, this subset of equilibria is most comparable to the behavioral model presented in the next section, which also obeys symmetry. I show that every symmetric, pure-strategy⁹ Bayes-Nash equilibrium predicts total pledges to meet the threshold exactly in certain state of nature. Provision occurs in some states

⁹Looking at the Volunteer's Dilemma, Weesie (1994) demonstrates a less-is-more prediction, by comparing mixed-strategy Nash equilibria under complete information to pure-strategy Bayes-Nash equilibria under incomplete information. As stated earlier, Goeree et al. (2005a) reject the mixed-strategy explanation for group size effects. This is why I focus exclusively on pure-strategy equilibria.

and not in others, so average provision rates across states range from zero and one.

To demonstrate this fact, it is first necessary to specify a parameterization of the game because the best-responses occur at discontinuities in payoffs that are non-differentiable. At these points, any infinitesimal reduction in pledges would strictly reduce the number of states in which the threshold is achieved, generating a sudden drop in expected payoffs. Payoff heterogeneity is generated by re-drawing each player's value, v_i , every round as either "high" or "low" with equal probability. Specifically, the low value is $v_l = 20$ while the high value is $v_h = 40$. Endowments are equalized to $w_i = 40$, which is helpful for comparison across states, but means that only a low type may contribute more than her value. The threshold is set equal to $T = nv_l$ in order to fix the average step return (SR) across group-size. Small groups have $n = 4$ members while large ones have $n = 8$. In this framework, it is weakly efficient to produce the public good across all states of the world; strictly so if there is at least one high type.

By restricting attention to the set of symmetric Bayes-Nash equilibria, strategy profiles simplify to an ordered pair, (c_l^*, c_h^*) , representing the type-contingent best responses of all players. I solve for the set of equilibria through numerically guess-and-check of every strategy profile in a finely discretized grid of the strategy space. Additional details regarding the numerical methods use here are found in

Appendix A. In these equilibria, total pledges perfectly match the threshold for certain states of nature.¹⁰ For example, the equilibria that exactly achieve the threshold in the state with k low types satisfy the equation below.

$$kc_l + (n - k)c_h = T \tag{1.2}$$

However, equilibria do not necessarily exist for every state, k , and only a subset of the strategy profiles that do satisfy these equations can support an equilibrium. The estimated equilibria are presented in Table 1.1. All equilibria have an average provision rate between 0 and 0.6875 for small groups and between 0 and 0.3633 for large groups. Hence, group-size influences the range of potential provision rates. If the coordination difficulty depends upon the number of players, then the realized provision rate will vary with group-size in ways that cannot be modeled using best response behavior.

¹⁰Attempts to verify these equilibria lead to systems of high order polynomial inequalities. Thus, verifying solutions would require a numerical implementation of Newton's method, which yields no more insight than the original grid search.

Table 1.1: Bayes-Nash Equilibria

(a) Groups of 4	
Set of Equilibria	Ex-ante Provision Rate
$c_l = 0$ and $c_h = 0$	0.00%
$c_l = 0$ and $c_h = 20$	6.25%
$c_l + 3c_h = 80$ for $0 < c_l < 12.84$	31.25%
$2c_l + 2c_h = 80$ for $13 \leq c_l \leq 18.33$	68.75%
(b) Groups of 8	
Set of Equilibria	Ex-ante Provision Rate
$c_l = 0$ and $c_h = 0$	0.00%
$c_l = 0$ and $c_h = 20$	0.39%
$c_l + 7c_h = 160$ for $0.4 \leq c_l \leq 7.4$	3.52%
$2c_l + 6c_h = 160$ for $0.2 \leq c_l \leq 11.9$	14.45%
$3c_l + 5c_h = 160$ for $8.5 \leq c_l \leq 15.5$	36.33%

These results are not meant to be an exhaustive characterization of the range of Bayesian Nash equilibria. Rather, it simply underscores the fact that (a) the conflicting-interest coordination problem remains under incomplete information, (b) the average provision rate is predicted to be less than one under perfect coordination, and (c) failure to coordinate will determine overall welfare. It is not possible to make more specific predictions across group-size or information treatments given the indeterminacy of these solution concepts. One might consider asserting additional restrictions that resolve the equilibrium selection problem, but what if individuals fail to stably coordinate? I argue in favor of using logit quantal response equilibrium, an intuitive model of coordination failure, for mak-

ing meaningful predictions across treatments and variables that were previously undetermined.

1.2.2 Quantal Response Equilibria

Given the difficulty of maintaining total contributions at the threshold, let alone attaining this goal in the first place, it is no surprise that experiments typically find noisy decision-making and frequent coordination failure. Accordingly, the pure-strategy equilibria that tacitly assume perfect coordination are of limited use in practice, regardless of whether theoretical solutions to the equilibrium selection problem can be obtained. It is natural to wonder whether the concept of mixed-strategy Nash equilibria can provide greater insight here.

To begin thinking about this issue, recall that there exists a set of dominated pure-strategy equilibria. Any mixed-strategy profile that nests all probability mass inside of this region is a Nash equilibrium too. This includes all manner of exotic mixing where, for example, a player may evenly distribute probability density along a compact subset of her choices, while placing the remaining probability upon a discrete mass point. Clearly, there will be infinitely many such equilibria, so mixed strategies do not help with the problem of indeterminacy. While other mixed-strategy equilibria may exist at or near the threshold, it is not clear how to solve for them. It stands to reason that the threshold would be missed with

positive probability in some cases, so mixing equilibria may lead to a reduction in the provision rate relative to pure-strategies. It strains credulity to argue that players are capable of mixing in such a sophisticated manner as to maintain the indifference of all other players, but that they are not just as easily capable of coordinating on a pure-strategy equilibrium at the threshold; especially when such mixing implies greater coordination failure and lower expected payoffs.

An alternative approach promoted by this research is to use the logit quantal response equilibrium (QRE) model originally proposed by McKelvey and Palfrey (1995). In this framework, players are no longer assumed capable of perfect best response to their incentives. Instead, each player chooses each action in some proportion to its respective benefit, reflecting the an imperfect attempt at best response. A degree of decision error permeates everyone's choices, and these errors have a weak feedback loop through expected payoffs. QRE accounts for all of this because it is an equilibrium model, where expected payoffs and quantal responses are mutually consistent and jointly determined.

To motivate this equilibrium concept, I refer to the dynamic model of noisy directional learning proposed by Anderson et al. (2004). In this model, players observe each other's choices in continuous time, and they regularly adjust their decisions over a continuous action space in the direction of increasing expected payoffs. While each player correctly calculates the directional change necessary

to increase her expected payoffs, she does not correctly anticipate the magnitude of the change required to perfectly best respond. Presumably, this is because she cannot accurately predict the choices of others, which is a sensible assumption in the context of multiple equilibria. Anderson et al. (2004) specify this decision error as a Weiner process, which admits a stochastic steady state with a logit specification. They name this steady state the “logit equilibrium”; although, it is simply a continuous analogue to logit QRE.¹¹

The authors further demonstrate that this logit equilibrium is locally stable under noisy directional learning dynamics for every potential game, and that such stability is global when equilibrium is unique. A potential game is a game for which there exists a potential function, $V(c_1, \dots, c_n)$, such that $\partial u_i / \partial c_i = \partial V / \partial c_i$ for every player i . One sufficient condition for the existence of a potential function is that payoffs may be written in the following form:

$$u_i(c_i, c_{-i}) = u(c_i, c_{-i}) + \theta_i(c_i) + \phi_i(c_{-i}) \quad (1.3)$$

Suppose for a moment that excess contributions are not refunded in the threshold game. If the threshold is achieved, each player’s payoffs are $u_i = v_i - c_i$, otherwise they are zero. It is clear that this is a potential game where $u(c_i, c_{-i}) = \phi_i(c_{-i}) = 0$

¹¹McKelvey and Palfrey (1995) only prove the existence of logit QRE for finite games. However, payoff functions in the threshold obey step-wise continuity; a sufficient condition for existence of the continuous analogue to logit QRE.

and $\theta_i(c_i) = v_i - c_i$. One could just as easily guess the potential function to be the sum of all payoffs $V(c_1, \dots, c_n) = \sum_i (v_i - c_i)$ since $\partial V / \partial c_i = \partial u_i / \partial c_i = -1$ for any player i . Hence, the principal branch of logit equilibrium exists and is locally stable under noisy directional learning dynamics in the absence of the proportional refund. This logic does not translate directly to the case where there is a refund, which is not a potential game. However, the reasoning underpinning this stability is simple, intuitive, and applicable to both scenarios. Whenever total contributions are below the threshold, everyone's expected payoffs are increasing in the direction of increased contributions, and vice versa when total pledges exceed the threshold. The proportional fund only dampens the magnitude individual costs when the threshold is exceeded, which is unlikely to substantively alter the forces of attraction responsible for local stability.

In any case, the analysis presented here presumes that such a stochastic steady state has already been reached. This equilibrium is implicitly defined by the following two conditions.

- i. A set of quantal response functions, $\{\sigma_i\}_{i=1}^n$, that monotonically transform expected payoffs into probability densities according to the logit rule:

$$\sigma_i(c_i) = \frac{\exp(\lambda \pi_i(c_i))}{\int_{s_i} \exp(\lambda \pi_i(s_i)) ds_i} \quad (1.4)$$

- ii. A set of expected payoff functions, $\{\pi_i\}_{i=1}^n$, that correctly account for the noisy behavior of others:

$$\pi_i(c_i) = u_i(c_i, \sigma_{-i}) \tag{1.5}$$

It is clear that QRE is unique for $\lambda = 0$, since every player must randomize uniformly over all actions. McKelvey and Palfrey (1995) prove that the uniqueness of QRE holds for every $\lambda \in [0, \lambda^*]$, where λ^* is a constant that is game-specific and sensitive to parameterization. The authors also prove the existence of a continuous correspondence of equilibria connecting uniform mixing at $\lambda = 0$ to a generically unique Nash equilibrium as $\lambda \rightarrow \infty$, and they call this correspondence the “principal branch”.¹² This is the only subset of logit QRE that have been applied to data in the literature. The justification for this includes the fact that (i) it is the only set of equilibria that exists for certain, (ii) it is the only set of equilibria that exist for $\lambda \in [0, \lambda^*]$, and (iii) it is the only set of equilibria with a known method for solving numerically. Thus, the principal branch nests utter

¹²In symmetric or near-symmetric games with multiple equilibria, such as when players share a homogeneous value for the public good, the principal branch, itself, may no longer be unique for $\lambda > \lambda^*$. In such cases, the principal branch intersects other sets of quantal response equilibria, so it is no longer possible to use QRE to avoid the indeterminacy of multiple Nash equilibria. Likewise, the selection of a Nash equilibrium as a limiting QRE becomes a meaningless exercise because the principal branch nests more than one Nash equilibrium in the limit. Fortunately, such bifurcations occur far away from the typical estimates of λ found with experimental data, so this theoretical limitation is rarely binding in practice.

randomization and mutually consistent best response as special cases in a model of mutual quantal response.

While the logit specification prevents closed-form solutions to the correspondence of fixed points along the principal branch, modern homotopy methods may be used to numerically solve the system in a step-by-step procedure. I employ a variant of the algorithm described in Turocy (2005) to numerically approximate the principal branch of logit QRE under complete information. The Harsanyi transformation allows me to recast the game of incomplete information as a game of imperfect information regarding “Nature”. Thus, the relevant solution concept under private information is the extensive-form analogue to QRE, known as “Agent QRE” (McKelvey and Palfrey, 1998).

I solve for these equilibria using the algorithm described in Turocy (2010). This is perfectly consistent with noisy directional learning, as Anderson et al. (2004) prove their results for so-called population games. Both of these algorithms are extensions of the code incorporated into The Gambit Project (McKelvey et al., 2013); a set of open-source tools for computational game theorists.¹³ These extensions are designed to exploit “mean-statistic” games where each player’s preferences only depend upon the sum of others’ choices. Exploiting this special feature makes it computationally feasible to solve games with more than two players.


¹³Many thanks to Theodore L. Turocy for generously sharing this code

1.2.3 Quantal Response Equilibrium Predictions

Once the quantal response functions are obtained by numerical methods, it is possible to make predictions about the mean and variance of pledges by low and high payoff players. It is also possible to construct the distribution of total contributions and make predictions about the mean and variance of this distribution. Moreover, one can measure the probability mass of pledged contributions above an individual's valuation for low types, or the probability mass of total contributions in excess of the threshold, which predict the over-contribution rate and the provision rate, respectively.¹⁴ Using the parameterization of the threshold game discussed in section 1.2.1, Figures 1.1 and 1.2 plot three points along the principal branch correspondence across all four experimental treatments for $\lambda \in \{0.2, 0.4, 0.6\}$. Figure 1.1 plots the individual quantal response functions by payoff-type, while Figure 1.2 presents the distribution of total contributions and the threshold. Notice that the predicted values of all variables of interest, including the mean and variance of individual and total contributions, as well as average over-contribution and provision rates, vary in a non-linear fashion across all treatments for any given value of λ .

¹⁴Note that these predictions are averaged across states, as many of the estimated moments in the experimental data must be averaged across states as well. For now, these means are constructed using the binomial probabilities that apply ex-ante. I re-weight these predictions using the frequencies found in the experiment for the predictions provided in section 1.3.3.


Figure 1.1: QRE: Individual Contributions



As a stochastic model of choice, the response functions of QRE may be used to construct a structural maximum likelihood estimator for the precision parameter, λ . While this value has the same units as payoffs, it does not have an absolute interpretation. Rather, it is a relative measure of responsiveness to expected payoffs, i.e. it is a relative measure of how well incentives predict choices. A note of

caution is appropriate here: λ is implicitly embedded inside of σ_{-i} , so the precision parameter is only a relative measure of responsiveness to the endogenously determined expected payoff functions, $\pi_i(\cdot)$. This is because the logit equilibrium is a closed model where quantal responses and expected payoffs are simultaneously determined. An increase in relative responsiveness may sharpen or dull incentives, so a rise in λ can potentially increase or decrease the mean or variance of individual or total pledged contributions. In certain cases, increased payoff responsiveness can *increase the noisiness* in decision-making for some individuals.

Figure 1.2: QRE: Total Contributions



Typically, economic models track the evolution of an endogenous parameter as a function of exogenously specified shocks. This model works in reverse, by tracking the evolution of endogenous shocks as a function of an exogenously varying parameter. As such, there is little *a priori* reason to expect the exogenous parameter to be constant across games or treatments. For example, McKelvey and

Palfrey (1995) and Rogers et al. (2009) find parameter estimates that vary significantly across games, while Sheremeta (2011) finds similar differences within four implementations of simultaneous contests. Allowing additional degrees of freedom for each treatment risks the criticism of over-fitting, while fixing one degree of freedom across treatments risks misspecification error. With no prior reason to impose the restriction, I allow the logit parameter to vary across treatments; although, this has been criticized as post-hoc and empirically vacuous (Haile et al., 2008).

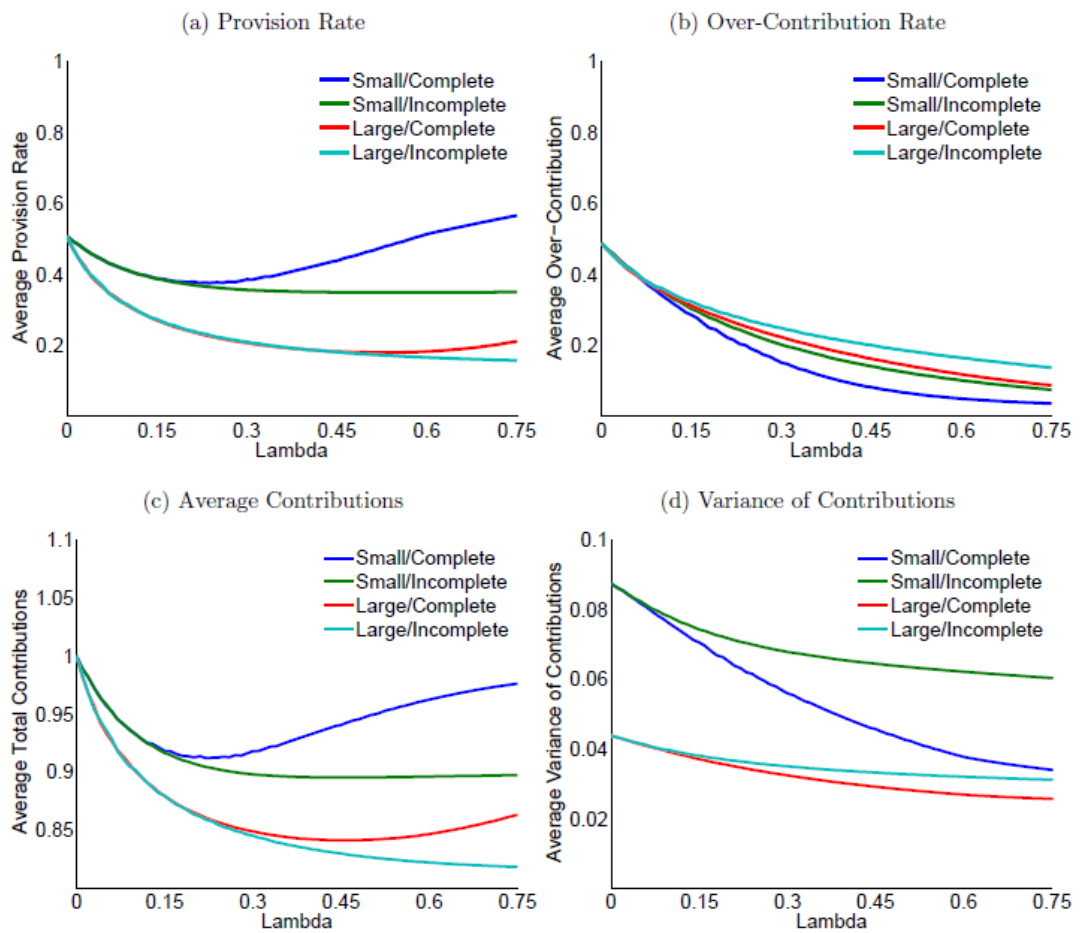
To preempt such doubts or concerns, I will first lay out some properties of the model found across a range of parameter values. The model makes a host of directional predictions for values of $\lambda \in [0, 0.75]$.¹⁵ These predictions correspond to some simple economic intuition of decision errors, and they provide the primary insights of the model; although, the specific magnitudes of such effects ultimately require estimation.

In response to the criticism from Haile et al. (2008) that QRE lacks empirical content, Goeree et al. (2005b) define a list of regularity conditions necessary for QRE to place testable restrictions upon the data: interiority, continuity, responsiveness, and monotonicity. All of these are satisfied by the logit specification

¹⁵It would be more satisfying to make claims regarding every $\lambda \in [0, \infty)$. Unfortunately, the bifurcations mentioned in section 1.2.2 make numerical calculations difficult when all players share the same payoff type under complete information. Nonetheless, the interval used here is sufficiently large as to nest all of the parameter estimates found in the data. I discuss this issue in greater detail towards the end of this section.

with independent preference shocks. Indeed, the directional predictions I present below are genuinely falsifiable restrictions on the data, with only one degree of freedom required to predict the level of eight separate variables. This reasoning is similar in spirit to the analysis presented by Goeree et al. (2005b) that yields testable hypotheses for the matching pennies game imposed by all regular quantal response equilibria.

Figure 1.3: Group-Level QRE Predictions



Figures 1.3 and 1.4 plot the group-level and individual-level predictions, respectively. Note that the group-level variables considered in Figure 1.3 are state-contingent. Hence, averages must be taken across all states to collapse these predictions into a single figure, and these averages correspond to the observed moments in the experimental data where states are rotating round-to-round. Figure 1.3(a) plots the average probability that total contributions reach the threshold, and Figure 1.3(b) plots the average probability that a low payoff player pledges more than her value to the public good. Figures 1.3(c) and 1.3(d) plot the average mean and variance of total contributions normalized by the threshold. This normalization is useful for making clean comparisons across group-size, where the threshold and the range of total contributions scale proportionally.

Several clear patterns emerge regarding the relevant treatment effects for any value of the precision parameter. First, the provision rate and average contributions decrease when information is reduced or the number of players is increased. Moreover, the effects of group-size appear to dominate the effects of information. For low values of λ , information effects are negligible, since players who do not respond to payoffs must necessarily be ignoring the information contained in their incentives. As precision increases, payoff uncertainty has a larger impact on outcomes; though, even large groups with complete information cannot catch up to small groups with incomplete information for the parameter values presented here.

Unsurprisingly, the average variance of contributions is decreasing in λ , implying that increased sensitivity to payoffs is associated with an overall reduction in noise. This finding is highly pronounced for complete information, though the dissipation of noise is noticeably slower under incomplete information. Lastly, the over-contribution rate of low payoff players decreases in tandem with this reduction in noisiness. Therefore, the over-contribution predicted by QRE represents a pure decision-error that vanishes as behavior approaches best response.

Figure 1.4: Individual-Level QRE Predictions




Figure 1.4 presents the findings at the individual level. Figures 1.4(a) and 1.4(c) plot the mean of individuals contributions of low and high payoff players, respectively, while Figures 1.4(b) and 1.4(d) do the same for the variance of individual contributions. While the standard deviation would be a more natural for interpreting the units of noisiness, I focus attention on the variance for the sake

of statistical tests. When examining the data, it will be necessary to statistically detect significant changes in variability, and the sample variance estimator has a well-known distribution for testing such hypotheses. Each player conditions her quantal response upon the full state of the world under complete information, so the mean and variance of individual contributions may vary state-by-state; while the opposite is true of incomplete information, by definition. The values presented here take averages of these distributional moments across all states of nature for this reason.

To begin, notice that average contributions start at twenty for both payoff types in Figures 1.4(a) and 1.4(c). This is because all players have the same choice set, $[0, 40]$, and uniform mixing implies a mean right in the center. For low payoff players, this means an average benefit of zero, so it is no surprise that their average contributions are decreasing as they become less sloppy in decision-making. Such a reduction in average contributions would significantly reduce expected payoffs via a reduced probability of success, if it were not for the fact that high types increase their contributions to compensate. Thus, I am left with the weak inequality that the average contribution of high payoff players is at least as large as that of low types. These changes are greater under complete information and for smaller groups, though such differences are only minor. Likewise, increased payoff responsiveness is associated with a reduction in individual noisiness in

Figures 1.4(b) and 1.4(d), and that these changes are more pronounced for smaller groups. For low payoff players, this reduction in variance is faster under complete information, but the difference across information treatments is imperceptible for high types.

Before proceeding to the experiments, it is important to say something about the limiting logit equilibria in the threshold game. Recall that there are uncountably many efficient Nash equilibria that meet the threshold, and there are also uncountably many dominated Nash equilibria that fall well below the threshold. It is quite possible that there are uncountably many logit quantal response equilibria for $\lambda > \lambda^*$, and all of these equilibria may even be a part of the principal branch. These equilibria may connect to the efficient outcomes in the limit, or they may connect to the dominated outcomes (or both). If the principal branch exclusively nests the efficient outcomes at the threshold, then we might expect certain properties of QRE to carry through in the limit. For example, the comparative static described above that average contributions be greater for high payoff players than low types may restrict the limiting equilibria to contributions (c_l, c_h) such that $c_l \leq c_h$.

On the other hand, all bets are off if the limiting QRE nests the dominated Nash equilibria, where high and low payoff players will both be weakly indifferent across all actions below their private valuations. This limits the usefulness of logit

QRE since the model cannot solve the underlying problem of equilibrium selection. Indeed, some parameterizations may admit multiple equilibria, regardless of whether behavior is assumed to follow best response or quantal response. Little more can be said on the matter, except that these issues of indeterminacy do not arise in the range of parameter values typically estimated using experimental data. Moreover, the thesis of this research is that quantal response behavior is a sensible and useful assumption for modeling coordination failure, so the limiting properties that lead back to pure best-response are of little interest.

1.3 The Experiment

The data collection took place within the Experimental and Behavioral Economics Laboratory at University of California, Santa Barbara. Sessions were run in the Fall of 2012 using the z-Tree software package (Fischbacher, 2007). Participants were randomly selected from a group of undergraduate and graduate students at the university who have registered interest to participate in social science experiments.

This experiment follows a two-by-two design, with small versus large groups and complete versus private information regarding payoffs. There are eight sessions, including two sessions for each treatment to control for session-specific effects. The public good is described in terms of a monetary prizes to avoid priming

social norms or altruistic motives common to public goods. Each session includes thirty rounds of just one treatment: the first ten rounds are for practice, while the latter twenty are potentially paid rounds. At the end of each session, one of the potentially paid rounds is chosen at random to determine payments for each subject. In addition, each participant receives a \$5 show-up fee, and an additional \$5 bonus, regardless of whether the threshold is reached. This bonus is used to penalize over-bidding by low types, and it could only be lost if a participant contributed more than her valuation when the threshold was reached.¹⁶ Each session lasts about an hour, with average total earnings per participant of \$15.04.

Group assignments remain static across the practice rounds and paid rounds, but change between sets of rounds. By fixing groups, participants have an opportunity to form beliefs regarding each others' strategies, which is a necessary prerequisite to successful coordination. Players are assigned to groups anonymously, and they are not able to communicate, so as to minimize tacit cooperation during the experiment or side-payments afterward. Players do observe the level total contributions¹⁷ at the end of each round, as some minimal information is required for subjects to learn in accordance with the underlying model.¹⁸ They are not

¹⁶This penalty was not included in the pilot session, creating a perverse incentive for low payoff players to over-contribute. Unsurprisingly, over-contribution was a major problem in this session.

¹⁷In this game, knowledge of total contributions is all the information required to pursue pure self-interest.

¹⁸The model requires learning about others' strategies in order for individuals to determine the direction of adjustment that will increase expected payoffs.

able to observe the individual actions of other group members, which reduces the likelihood of using punishment to enforce higher levels of cooperation. This also suppresses the information required to act on distributional preferences, or to punish individual behavior that is considered unfair.

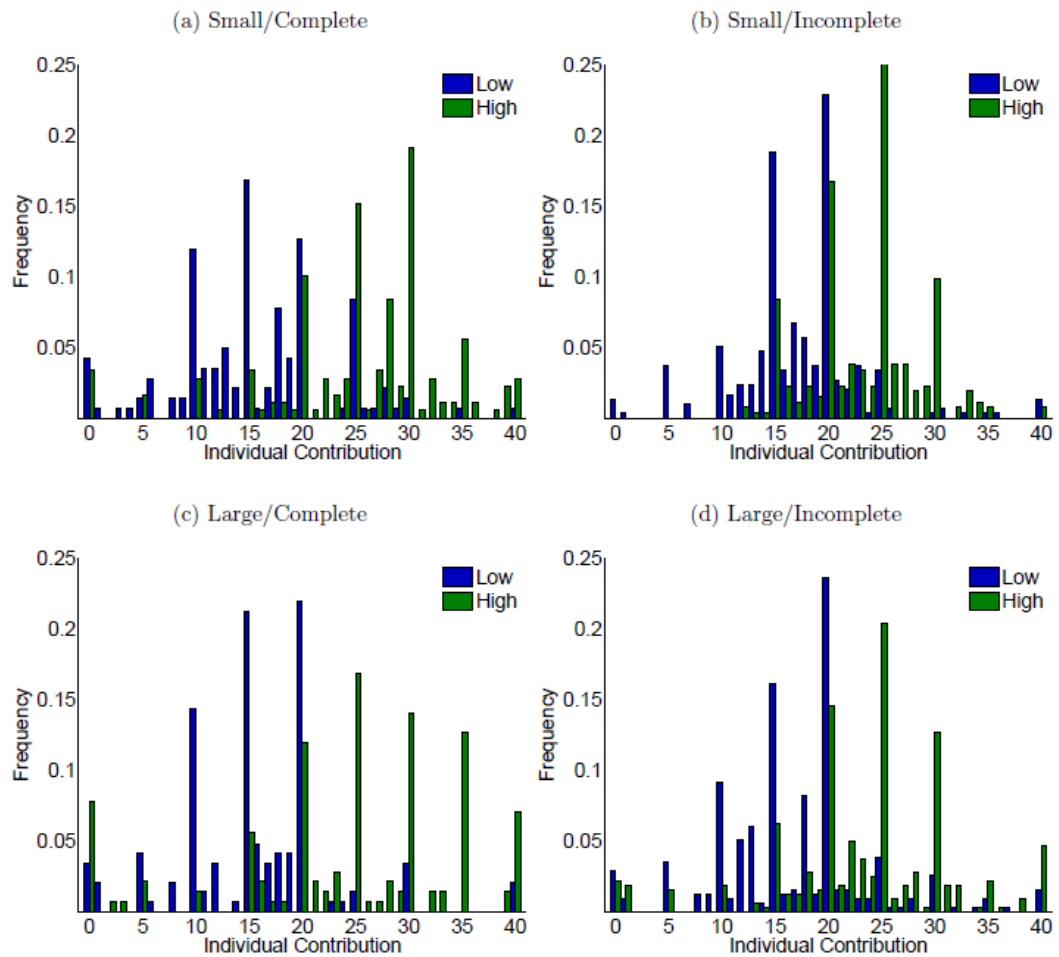
Small groups have $n = 4$ players, while large groups have $n = 8$ players. Payoff types are re-drawn independently and with equal probability in each round. Under complete information, each player knows the exact number of high and low types in her group each round, while under incomplete information, each player only knows her own type each round, as well as the ex-ante distribution of others' types across rounds. Each player's endowment is $w_i = 40$, the low value is $v_l = 20$, the high value is $v_h = 40$, and the threshold is $T = nv_l$. This ensures a constant average step return across group-size, while also guaranteeing the efficiency of provision in each state.

1.3.1 Analysis Across Individuals

On an individual level, observations are defined as a player's pledge in a given round. Each session of small groups includes four groups of four players, while each session of large groups includes two groups of eight players. Therefore, each session has sixteen subjects in total, with each subject playing twenty potentially paid rounds. Because each treatment receives two sessions, the sample size is

$N_i = 640$ per treatment. Unfortunately, one of the sessions for small groups with incomplete information did not have sufficient attendance for four groups. This session only had three groups of four players, for twelve subjects total. Therefore, this treatment had a sample size of only $N_i = 560$.

Figure 1.5: Histograms: Individual Contributions

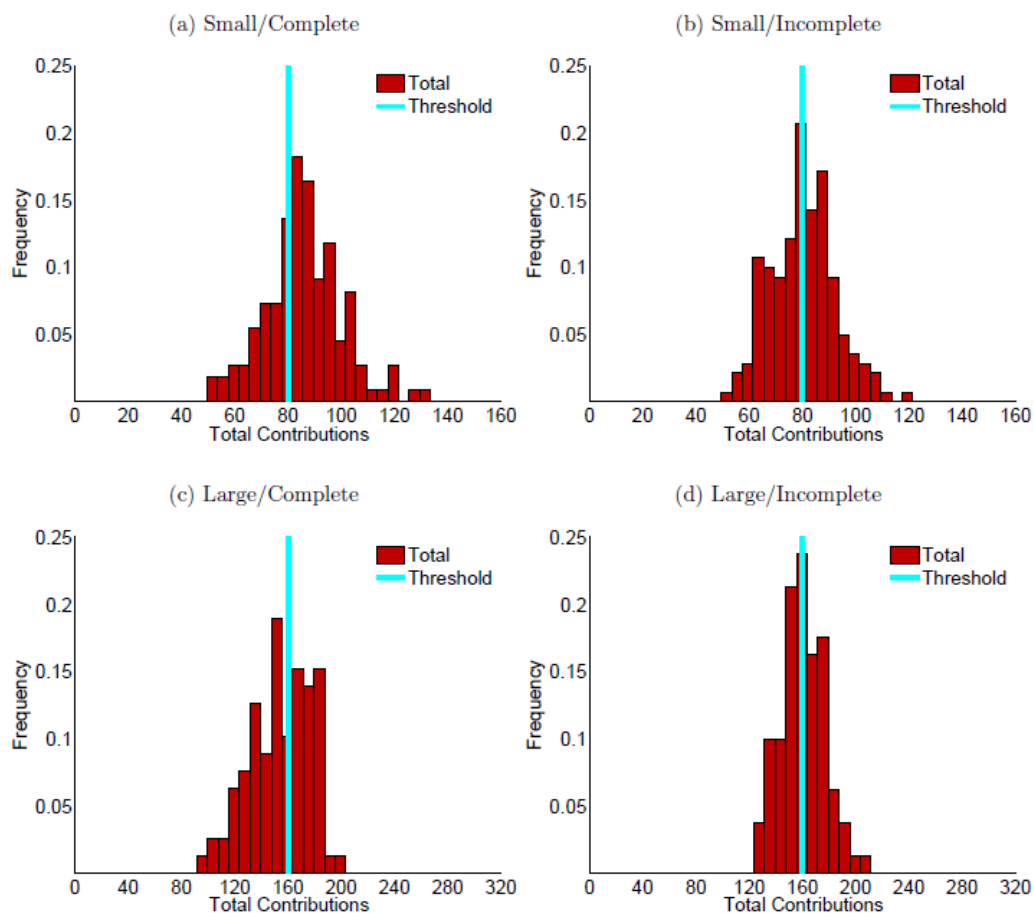


One pronounced feature of the individual data is the clustering of observations on integers, as can be seen in Figure 1.5. In particular, a majority of observations occur at multiples of five, suggesting that participants make their decisions as though they faced a discrete-choice problem. For this reason, estimates using a fine approximation of the continuous logit QRE may be biased. Fortunately, this clustering of data is not pronounced at the aggregate level, as can be seen in the histograms of total contributions presented in Figure 1.6. Two transformations of the data that may overcome this bias include (1) using a more discrete approximation of the model while sorting the data individual contributions into coarse histogram bins or (2) using total contributions instead of individual contributions to construct alternative maximum likelihood estimators.¹⁹

¹⁹Despite discrete approximation of the strategy space, the choice probabilities are still continuous. Therefore, key predictions such as the mean of contributions or the provision rate are hardly sensitive the coarseness of the discrete approximation. Coarseness should not change the model's predictions, only the potential estimate of λ .

Figure 1.6: Histograms: Total Contributions

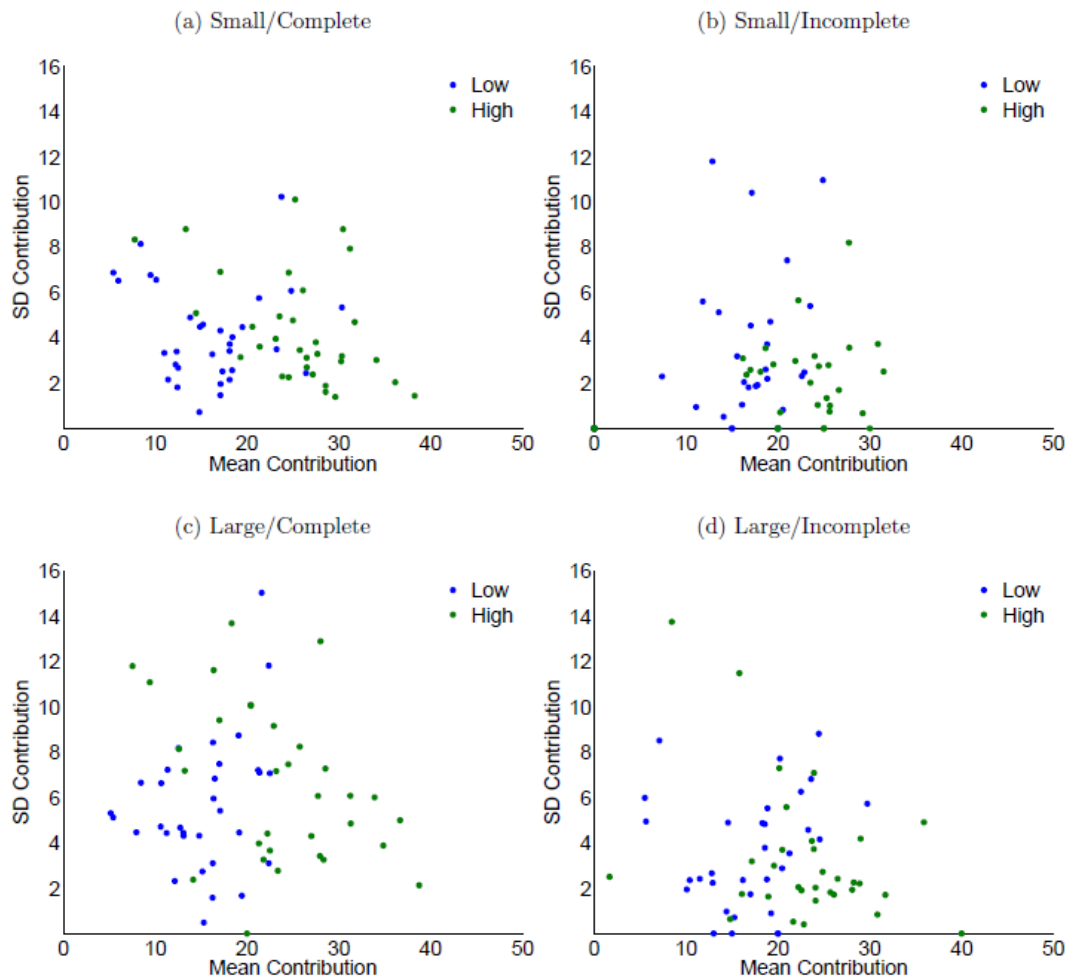
Figure 4: Histograms: Total Contributions



Interestingly, the overwhelming majority of individuals exhibited some variation in pledges across rounds, as can be seen in Figure 1.7. Under complete information, no players made identical pledges in every round as both types; although, one player made identical pledges as a high type with large groups. Under private information, four players made identical pledges as low types for both group-sizes.

For high types, five players showed no variation in small groups, while only one player admitted this behavior in large groups. Overall, only two players showed no variation in pledges across both payoff types, and both of these subjects were in small groups with incomplete information. Put bluntly, 98.4% of all participants appeared to be mixing.

Figure 1.7: Mean/SD Individual Contributions



Obviously, this does not prove that players are imperfect best-responders. Players may try to play a mixed-strategy equilibrium, or they may best respond to out-of-equilibrium beliefs that are updated across rounds. Nonetheless, variation in individual choice is a necessary, if not sufficient, condition for the applicability of noisy best response models.


1.3.2 Analysis Across Groups

At the group level, an observation is defined as the total contributions in a given round. To make the data comparable, I scale these totals by the threshold. With sixteen subjects per session, there are only two groups of eight per session, and two sessions per treatment. Given twenty potentially paid rounds per treatment, that leaves $N_g = 80$ group-level observations for each treatment with large groups. For small groups, the sample size varies by information treatment. Under complete information, there were four groups of four in each session, leading to a sample size of $N_g = 160$, while under incomplete information, one session had only three groups of four, implying a sample size of $N_g = 140$. All of these samples are large enough to detect variations in the distribution of total contributions for this experiment.

Indeed, there is substantial variation contributions across rounds, and there is little evidence of convergence.²⁰ Figure 1.8 graphs total contributions relative to the threshold, by group, for each treatment. The means and variances are stationary across rounds, which is consistent with the idea of a stochastic steady state. The overall differences in variance across treatments (recorded in Table 1.4) are visually apparent as well: specifically, that increased group-size reduces the variance of total relative contributions, while increasing information has the opposite effect. However, total contributions should be independent, as well as identically distributed, in order for QRE to be a useful “as if” model of aggregate behavior.

²⁰This is not surprising, as payoff types are re-drawn each round. If payoffs were static, then one might expect convergence to the threshold. However, it would not be possible to simulate incomplete information without randomized group re-matching.

Figure 1.8: Total Contributions by Group across Rounds



There is little evidence of persistence in total contributions across rounds. Table 1.2 presents the estimates of a first-order auto-regressive process in total relative contributions across each group in each treatment. Of the twenty-seven groups included in this experiment, only three present serial correlation at the 10% significance level. Unfortunately, each group experienced only twenty potentially

paid rounds, so there is hardly enough statistical power behind these results to be certain. Therefore, I also estimate a model with data pooled across groups.²¹ Contributions do exhibit persistence across rounds for small groups with complete information. Notice that group 8 had significant auto-correlation and average contributions well below the threshold. Restricting the sample to groups 1-7 removes any sign serial correlation, suggesting that the effect is concentrated solely within group 8. For every other treatment, there appears to be no serial correlation in total contributions across rounds.

Table 1.2: Serial Correlation in Total Contributions

(a) $TC_t = \rho_0 + \rho_1 TC_{t-1} + \varepsilon_t$

Case	Value	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 7	Group 8	All
Small	$\hat{\rho}_1$	-0.211 (0.228)	-0.448* (0.233)	-0.140 (0.238)	-0.311 (0.225)	0.039 (0.238)	-0.094 (0.256)	0.271 (0.221)	-0.490* (0.244)	0.181** (0.078)
Comp	$\hat{\rho}_0$	104.712 (3.098)	105.386 (2.703)	103.259 (1.897)	99.060 (1.484)	108.537 (4.469)	112.520 (2.951)	122.760 (5.535)	90.008 (1.997)	106.309 (1.664)
Small	$\hat{\rho}_1$	-0.106 (0.201)	-0.022 (0.242)	-0.039 (0.272)	0.022 (0.230)	-0.353 (0.219)	0.323 (0.218)	0.138 (0.251)	-	0.131 (0.084)
Inc	$\hat{\rho}_0$	99.177 (2.774)	110.395 (2.964)	106.431 (3.850)	105.393 (3.715)	92.524 (1.955)	103.233 (3.549)	88.728 (3.853)	-	100.443 (1.469)
Large	$\hat{\rho}_1$	-0.229 (0.250)	-0.173 (0.239)	-0.229 (0.235)	-0.428* (0.216)	-	-	-	-	0.056 (0.114)
Comp	$\hat{\rho}_0$	99.884 (2.360)	87.047 (2.851)	94.906 (2.633)	105.842 (1.120)	-	-	-	-	96.736 (1.670)
Large	$\hat{\rho}_1$	0.048 (0.272)	0.092 (0.242)	-0.064 (0.237)	-0.130 (0.217)	-	-	-	-	0.067 (0.114)
Inc	$\hat{\rho}_0$	99.108 (3.075)	95.887 (2.495)	103.134 (1.520)	104.792 (1.851)	-	-	-	-	100.275 (1.235)

Null hypotheses are $H_0 : \hat{\rho}_1 = 0$ or $\hat{\rho}_0 = 100$

²¹In order to combine the data, the last round of one group is followed by the first round of another group. Clearly, outcomes are independent across groups. Falsely encoding one in every twenty observations as potentially serially correlated, when they are certainly independent, introduces a small amount of attenuation bias in the estimates of serial correlation. However, the increase in sample size from 20 to 80-160 observations profoundly increases the likelihood of detecting any serial correlation.

1.3.3 Quantal Response Equilibrium Results

An estimate the logit parameter is required to make predictions with QRE. I estimate one parameter value per experimental treatment. For incomplete information, only two quantal response functions are required to construct the likelihood function; one for each payoff-type. The complete information case is a bit more complex. In the experiment, values are re-drawn every round, so each round may have a different number of high types and low types, i.e. a different state of the world. Presumably, players condition their choices on the state of the world under complete information, so the quantal response functions are state-dependent. The complete information treatment requires eight (sixteen) separate quantal response functions in order to construct the likelihood function for small (large) groups.

Table 1.3: Maximum Likelihood Estimates (all 20 rounds)

(a) $\hat{\lambda}_{mle}$ and se_{mle}		
Group-size	Complete	Incomplete
Small	0.320 (0.014)	0.523 (0.030)
Large	0.440 (0.027)	0.557 (0.045)

Table 1.3 presents the parameter estimates for each treatment, along with standard errors. It is immediately apparent that $0 < \lambda < \infty$ for every treatment. Therefore, players are not mixing uniformly, nor are they perfectly coordinating

on a Nash equilibrium. The only two estimates of λ that are not significantly different are the two incomplete information treatments. The fact these estimates vary across treatments is not terribly surprising, since the parameter has very little intrinsic interpretation across these resulting non-trivial transformations of payoffs.

Using these parameter values, it is now possible to make specific predictions using the QRE, and to compare these predictions with observed behavior. While the mean predictions of the QRE model plotted in Figures 1.1 and 1.2 use binomial probabilities across states, I use the realized frequencies of each state within the experiment to make average predictions in Tables 1.4-1.7. That way, any gap between the model's predictions and the measured variables reflects true errors in prediction; not a gap between the ex-ante and ex-post distribution of states. Tables 1.4 and 1.5 present the group-level comparisons, while Tables 1.6 and 1.7 do the same for individuals. Tables 1.4 and 1.6 present information in levels, while Tables 1.5 and 1.7 focus on changes (comparative statics).

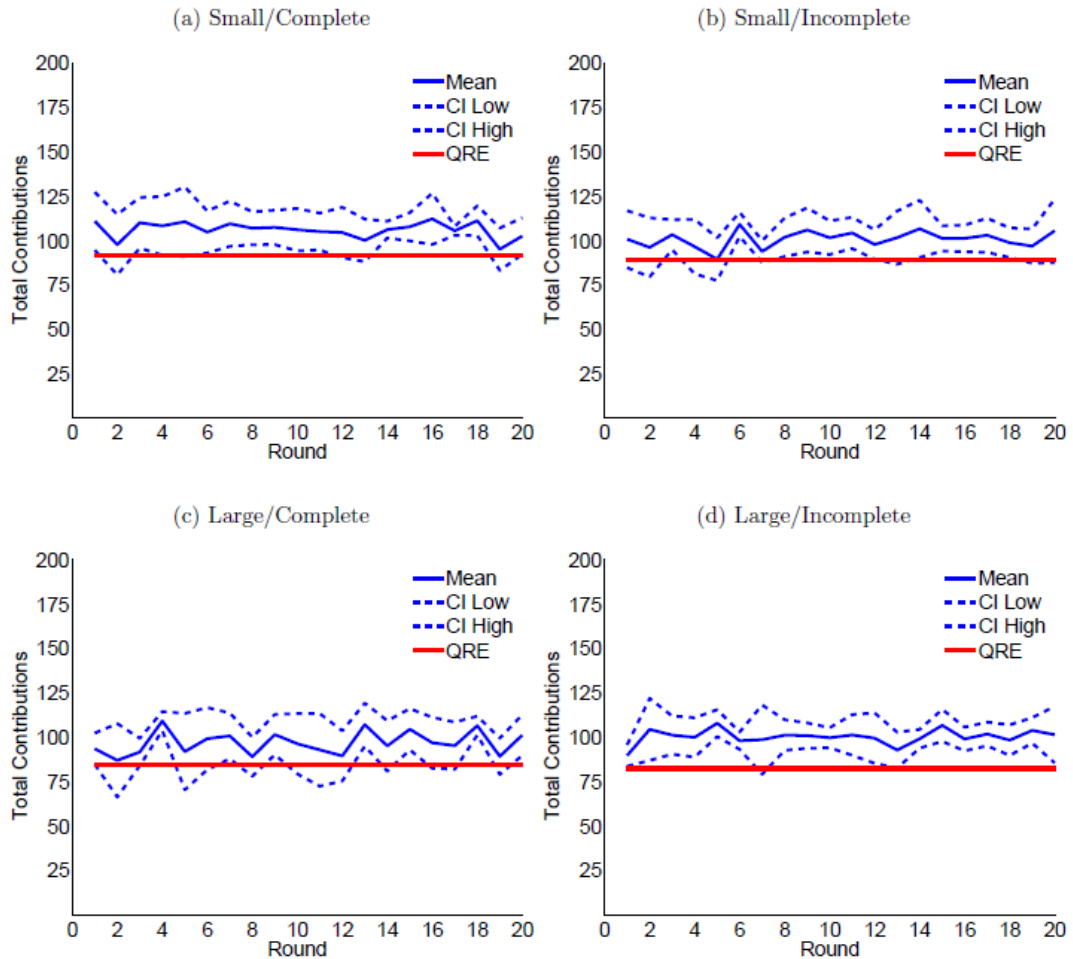
Table 1.4: Estimated and Predicted Levels for Groups

Variable	Estimate Prediction	Small Complete	Small Incomplete	Large Complete	Large Incomplete
Provision	Data	0.675 (0.037)	0.586 (0.042)	0.425 (0.055)	0.525 (0.056)
Rate	QRE	0.405	0.320	0.180	0.174
Total	Data	1.061 (0.014)	1.008 (0.013)	0.968 (0.016)	1.002 (0.011)
Contributions	QRE	0.920	0.895	0.841	0.823
Variance	Data	0.006 (0.030)	0.024 (0.005)	0.019 (0.005)	0.010 (0.003)
of Total	QRE	0.053	0.063	0.029	0.032
Over-contribution	Data	0.144 (0.028)	0.141 (0.029)	0.133 (0.038)	0.133 (0.038)
Rate	QRE	0.156	0.147	0.190	0.202

It is clear from Table 1.4 and Figure 1.9 that the QRE systematically under-predicts total relative contributions and average provision rates. Likewise, the model over-predicts the variance of total contribution. Both of these findings are consistent with the notion that λ is biased downward by a false assumption of homogeneity (Golman, 2011). Nonetheless, there is one dimension on which QRE truly shines. Several observations include over-bidding by low types. Regardless of the reason for this behavior, the QRE predicts the frequency of over-contribution quite well. This stands in stark contrast to Nash predictions, which do not allow a player to over-contribute in equilibrium.²²

²²Technically, the Nash equilibria do not allow an individual to bid more than her valuation in equilibrium, unless total pledges are so low that players find themselves in the dominated region of the strategy space. Contributions were consistently close to the threshold throughout this experiment, so this explanation of over-contribution lacks credibility.

Figure 1.9: Total Contributions by Group across Rounds



Looking at differences as opposed to levels (Table 1.5), the evidence is a bit more nuanced. Increasing information raises average contributions for small groups by 5.3% relative to the threshold ($p < 0.01$), but QRE predicts an only half of that increase at 2.5%. Meanwhile, increasing group-size reduces the provision rate by 25% points ($p < 0.01$) and reduces relative contributions by 9.3%

($p < 0.01$) under complete information, while QRE predicts very similar decreases of 22.5% points and 7.9%, respectively. Lastly, increasing group-size lowers the standard deviation of average contributions for large groups by 5.2% relative to the threshold ($p = 0.01$), while the logit equilibrium predicts a reduction of 7.2%. These are the only four group-level comparative statics with significance at or below the 5% level, and the sign of each effect is correctly predicted by QRE.

Table 1.5: Information and Group-size Effects for Groups

Variable	Estimate Prediction	Info Effect Small	Info Effect Large	Size Effect Complete	Size Effect Incomplete
Provision	Data	0.089 (0.056)	-0.100 (0.076)	-0.250*** (0.067)	-0.061 (0.070)
Rate	QRE	0.085	0.006	-0.225	-0.146
Total	Data	0.053*** (0.019)	-0.034* (0.019)	-0.093*** (0.021)	-0.006 (0.017)
Contributions	QRE	0.025	0.018	-0.079	-0.072
Variance	Data	0.007 (0.008)	0.009* (0.005)	-0.011 (0.007)	-0.013** (0.005)
of Total	QRE	-0.010	-0.00.	-0.024	-0.031
Over-contribution	Data	0.003 (0.040)	0.000 (0.054)	-0.011 (0.047)	-0.008 (0.048)
Rate	QRE	0.009	-0.012	0.034	0.055

Info Effect = Complete - Incomplete, Group-size Effect = Large - Small

For large groups, there is some evidence that increasing information reduces average total contribution by 3.4% relative to the threshold ($p = 0.08$), in addition to increasing the standard deviation of by 3.7% ($p = 0.09$). At first, this appears to be an example of the “less-is-more” effect of information under bounded rationality described by Gigerenzer and Goldstein (1996), wherein a reduction of information leads to decisions that increase expected payoffs. Upon close inspec-

tion of the data, this finding is instead driven by outlier performance from group 2 in the treatment with large groups and complete information, dragging down the performance of the complete information treatment. This particular group only managed to achieve the threshold in 20% of the potentially paid rounds, and had an average of total contributions 15% below the threshold. Because the large-group treatments only contain four groups each, one outlier group can have significant impact on overall estimates.

A summary of individual contributions can be found in Table 1.6. Unsurprisingly, I find that high value players tend to pledge more, on average, than low value players. The QRE is able to replicate this basic feature. With the exception of average contributions by high value players in small groups, the model again underestimates the average level and overestimates the variance of individual pledges across all treatments and payoff types. Unfortunately, this gap in prediction reinforces a key short-coming of this work: the need for heterogeneity in the logit parameter.

Table 1.6: Estimated and Predicted Levels for Individuals

Variable	Value	Estimate Prediction	Small Complete	Small Incomplete	Large Complete	Large Incomplete
Individual Mean	Low	Data	16.179 (0.449)	17.169 (0.445)	15.151 (0.431)	16.753 (0.431)
	Low	QRE	11.309	11.459	11.922	12.220
	High	Data	25.461 (0.411)	23.531 (0.473)	23.498 (0.426)	23.156 (0.426)
	High	QRE	24.620	24.358	21.424	20.735
Individual Variance	Low	Data	49.939 (6.238)	35.924 (6.185)	57.210 (5.987)	51.683 (5.987)
	Low	QRE	58.571	52.601	67.285	69.633
	High	Data	63.069 (5.714)	26.022 (6.573)	108.168 (5.931)	67.296 (5.931)
	High	QRE	91.459	66.596	110.157	100.822

Table 1.7 presents the differences that exist at the individual-level. The average contributions of high types exceeds that of low types in every treatment. Also, high value players admit a greater variance in contributions, but only in small groups under complete information. More information increases the average pledges of high types in small groups, while it decreases the average pledges of low types in large groups. In addition, increasing information increases the variance of contributions by high types, regardless of group-size. Increasing group-size reduces the average pledges of high types under complete information, but it increases the variance of high types' contributions under both information complete and incomplete information. All twelve of these comparison are significant below the 1% level, and QRE correctly predicts the sign of each one; though it certainly has trouble predicting the magnitudes of these differences.

Table 1.7: Information, Group-size, and Payoff-Type Effects for Individuals

		(a) Effect on Mean			
Variable		Data Low	Data High	QRE Low	QRE High
Information	Small	-0.990 (0.632)	1.930*** (0.626)	-0.15	0.262
	Effect	-1.602*** (0.609)	0.242 (0.603)	-0.298	0.689
		Low	High	Low	High
Group-size	Complete	-1.029* (0.622)	-1.963*** (0.592)	0.613	-3.196
	Effect	-0.416 (0.619)	-0.274 (0.637)	0.761	-3.623
		Complete	Incomplete	Complete	Incomplete
Payoff-type	Small	9.281*** (0.608)	6.361*** (0.649)	13.311	12.899
	Effect	8.347*** (0.606)	6.503*** (0.606)	9.502	8.515

		(b) Effect on Variance			
Variable		Data Low	Data High	QRE Low	QRE High
Information	Small	14.015 (8.785)	37.048*** (8.710)	5.97	24.863
	Effect	5.527 (8.467)	40.873*** (8.389)	-2.348	9.335
		Low	High	Low	High
Group-size	Complete	7.271 (8.646)	45.099*** (8.236)	8.714	18.698
	Effect	15.758* (8.608)	41.273*** (8.854)	17.032	34.226
		Complete	Incomplete	Complete	Incomplete
Payoff-type	Small	13.131 (8.460)	-9.902 (9.026)	32.888	13.995
	Effect	50.959*** (8.428)	15.613* (8.428)	42.872	31.189

Info Effect = Complete - Incomplete, Group-size Effect = Large - Small, Payoff-type Effect = High - Low

1.4 Discussion

Many lab experiments find anomalous behavior regarding the provision of public goods. For example, the Nash equilibrium level of individual contributions is zero in the canonical linear public good game, while the socially optimal outcome is for everyone to contribute as much as possible. As has been documented by an extensive literature (Davis and Holt, 1993; Kagel and Roth, 1997), observed contributions are typically greater than the free-rider level but less than the optimal level. Two theories, in particular, have been promoted to explain these stylized facts: (i) decision error and ii) altruism. It is not clear how to separately identify each motivation in the linear public goods game, as any strictly positive contribution below one's own endowment is rationalizable by either method.

Using a novel experimental design, Andreoni (1995) made significant headway on this identification problem.²³ Varying information regarding the relative rank of players' contributions, and in one treatment transforming payoffs into a zero-sum game over rank, the author isolates a lower bound on the fraction of contributions that are plausibly pure "kindness". Three quarters of all contributions were positive, about half of which are attributed to altruism. In follow up work, Andreoni and Miller (2002) conduct an experiment to test whether observed

²³The author does discuss the potential interaction of altruism and confusion, as notions of fairness often give rise to multiple equilibria, which may further exacerbate confusion.

giving in various parameterizations of the dictator game could also be rationalized by standard, quasi-concave preferences over payoffs of self and others. They find that 47% of subjects may be exactly characterized as either purely self-interested, utilitarian, or Leontief (perfect equity). The remaining subjects were grouped into these three categories by minimizing the distance from their choice to the respective predictions. Assuming CES utility over payoffs, the elasticity of substitution and payoff weights were estimated for each remaining group. The results strongly support the hypothesis that unobserved, heterogeneous altruistic preferences explain a great deal of observed giving.

Yet, there always remains room for decision error as well. Houser and Kurzban (2002) replicate the experimental design from Andreoni (1995), except they pit subjects against computers in one treatment to remove any last vestiges of social norms, reciprocity, or fairness considerations. They also estimate that half of positive contributions are altruistic and half are mistakes, perfectly in line with the results of Andreoni (1995). Palfrey and Prisbrey (1997) test two competing notions of kindness in a public goods game: pure altruism (care about others' payoffs) and warm-glow (enjoy the act of giving). The experiment rotates the private return to contributions to the public good to each player in each round, which also rotates the maximum total benefit to the public good. This should have large effects under pure altruism but no effect on warm-glow. The authors find

strong support for warm-glow but little evidence of pure altruism.²⁴ It would seem that the appropriate method to model altruism will always be context-specific, and it will depend on the experimental design and the relative importance of errors.

One unsatisfactory aspect of these previous papers is that they lack a formal notion of confusion, thus rendering the term as a catch-all for any non-Nash, non-altruistic behavior. By contrast, Anderson et al. (1998) propose using the logit equilibrium to model decision error by subjects in linear and quadratic public goods games. Contributions appear to increase stochastically with the marginal value of the public good, representing the kind of magnitude effects that are best described by models of stochastic choice. However, adding pure altruism and warm-glow substantially increases the fit of the logit equilibrium model to the data. Offerman et al. (1998) actually use QRE to model decision errors in contributions to the step return game (a discrete variant of the threshold game). The authors find the model's predictions useful, though the results favor a naive Bayesian model of stochastic choice, in which individuals' expectations of the others' contributions are overly optimistic. Unfortunately, this paper does not incorporate altruism into utility when comparing these two models.

²⁴Perhaps rotating payoffs across players and rounds complicates the environment which, in turn, exacerbates confusion. If this is true, then decision error will play an important role in the experiment conducted for this research.

The takeaway from all of this is that both altruism and decision error play a prominent role in explaining individual behavior in any public goods game. The debate is not a matter of one over the other, but simply a question of degree. Therefore, it seems sensible to account for both features explicitly. Of course, the validity of these estimates will depend upon the functional forms. It is common practice to use the simple specifications such as a linear function over payoffs.

$$u_i = \pi_i + \alpha\pi_{-i} \text{ for } \alpha \geq 0 \tag{1.6}$$

These preferences (perfect substitutes) lead to corner solutions that unambiguously bias contributions upwards in the threshold game. If $\alpha = 0$, the model simplifies to pure self-interest. For $0 < \alpha < 1$, each player places greater weight upon a unit increase her own payoff than for others. A player would be willing to contribute more than her valuation to the public good if she is pivotal in reaching the threshold, so long as the discounted benefit to others exceeds the potential cost to self. Yet, assuming one is already at the threshold, there would be no willingness to sacrifice an additional unit of payoff in order to redistribute payoffs to others through the over-contribution refund. In the case that $\alpha = 1$, payoff preferences represent pure utilitarianism. In the threshold game, any set of pledges greater than or equal to the threshold will be strictly equally preferred

to any set of pledges that fail, and players are indifferent between every set of pledges greater than or equal to the threshold, biasing individual contributions upwards and implying that total contributions should lie above the threshold. In the extreme case where $\alpha > 1$, players have a strict preference for maximizing the benefit of others' at the expense of self, without bound, which seems unlikely in most circumstances.

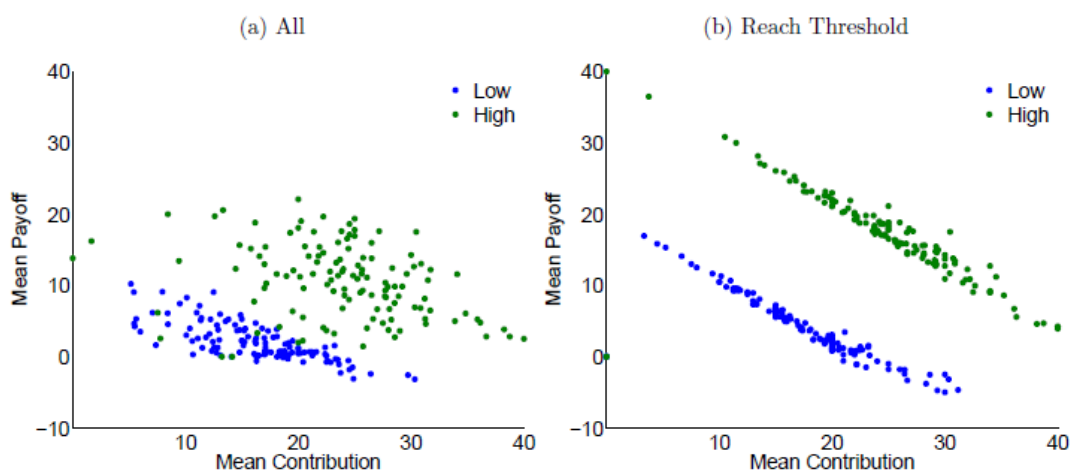
When considering warm-glow, an equally simply linear form is often sufficient:

$$u_i = \pi_i + \gamma c_i \text{ for } \gamma \geq 0 \tag{1.7}$$

Given the discontinuous structure of payoffs in the threshold game, it is best to consider the effect of warm-glow in two cases. If the threshold is not achieved, then utility is simply the warm-glow component $u_i = \gamma c_i$, which pushes incentives to contribute up towards the threshold, and away from the inefficient Nash equilibria with very little contribution. If the threshold is reached, then warm-glow increases utility while the reduction in monetary payoff reduces utility. How these forces net out is highly complex, in theory, given the non-linear nature of the refunded over-contributions. However, in the experimental data the refund appeared to have a relatively negligent effect on realized payoffs, as can be seen in Figure 1.10 when restricting data to when the threshold was at least achieved. For the sake of

parsimony, I will briefly assume the refund is off so that the combined utility from monetary payments and warm glow will be $u_i = v_i - (1 - \gamma)c_i$ when the threshold is reached.

Figure 1.10: Mean Individual Payoffs by Contribution



If $\gamma = 1$, then warm-glow cancels out the monetary loss of further contributions, making one indifferent between any set of contributions that reaches or surpasses the threshold, just as was the case for pure utilitarianism with $\alpha = 1$. For $\gamma < 1$ the warm-glow only partially offsets the monetary losses from further contribution, which means that only contribution vectors that perfectly achieve the threshold will be equilibria; although, these equilibria may include individuals contributing more than their valuations in some cases. This prediction coincides with the altruistic preferences with $\alpha < 1$. For $\gamma > 1$ the utility from warm-glow

swamps the monetary losses from contributing, suggesting that anyone should contribute her whole endowment. This does not seem like a useful description in most circumstances.

Another standard model of preferences over payoffs regards fairness. Perhaps individuals prefer a specific level or share of total benefits relative to others. Such ideas are captured by perfect complements preferences using a Leontief utility function over payoffs:

$$u_i = \min \{ \pi_i, \beta \pi_{-i} \} \quad \text{for } \beta \geq -1/2 \quad (1.8)$$

A Rawlsian would put equal weight upon each individual. That corresponds to a value of $\beta = 1/(n - 1)$, so that for every one unit increase in π_i there is a corresponding $n - 1$ unit increase in π_{-i} . For the threshold game, the total amount of net benefit is a fixed value; though, it is state-dependent and, therefore. Therefore, Leontief social preferences are equivalent to Cobb-Douglas social preferences in this game. If $\beta > 1/(n - 1)$, an individual prefers more than an equal share of the gains, and vice versa for lower values of β . In fact, preferences converge to pure self-interest as $\beta \rightarrow \infty$, and concern for others at the expense of self as $\beta \rightarrow -1/2$. For values in between, each player admits a desired share of $\beta/(1 + \beta)$.

Therefore, the full range of outcomes in which total contributions meet or exceed the threshold are rationalizable according to some parameter value.

There are a few complications in interpreting fairness with this data. First, notice that the total amount of net benefit from providing the public good is state dependent. Hence, a preference for fixed share of benefits is not only going to require differing contributions by payoff type, but also by state of the world. This added complexity may decrease the relative prominence of fairness concerns relative to decision errors. Likewise, there is some natural ambiguity as to whether equal net gains is the only salient notion of equity. For example, it may be that subjects wish to equalize contributions instead. In this experiment, every individual would pledge 20. The threshold would be met, exactly, in every possible state of nature. Low types will always receive a net gain of zero, while high types evenly split the gains. Further alternatives include benefit or cost shares that are distributed in proportion to private valuations of the public good, which could be represented using specific parameter values of Leontief utility. All of this the natural ambiguity regarding fairness, which is certainly common in real-world applications, is reason to expect a greater role for confusion.

As this experiment is designed, it is not particularly easy to separate out different types and parameterizations of altruism and fairness. This is due in part to the overlapping predictions of large sets of individual contributions at or

exceeding the threshold. It is also due to the natural ambiguities of the very notion of fairness in this environment with asymmetric and changing private values to the public good. Thankfully, pioneering work by Offerman and Sonnemans (1996) addresses many of these issues directly in the step return public goods game. They independently measure and categorize participants' social preferences, and they find evidence in favor of warm-glow altruism. Yet, they also find that individuals occasionally, but systematically, violate their own alleged preferences. To quote from their findings, "Although the present study clearly indicates that individuals do vary in their preferences regarding outcomes for others, it turns out to be very difficult to rationalise all choices. Some decisions simply seem to be errors from an economic point of view." Given the multiplicity of Nash equilibria under self-interest, altruism, warm-glow, and fairness, it should be no surprise that subjects are confused.

Evidently, what they are confused about are the ever-changing pledges of other group members. They do not appear to improve either, as there is no trend in total contributions or provision rates across rounds. The threshold is exactly achieved in only 4% of all potentially paid group-rounds, constituting 17 instances out of a total 460 group-rounds in the experiment. Clearly, it is difficult to coordinate total contributions directly onto the target threshold. Only in one of these cases did the group manage to repeat this achievement in a subsequent round; although, the

threshold was exceeded in fifteen cases and missed in two. Overall, there appears to be a bias in achieving provision, as one would expect with altruism. About 54% of total contributions exceed the threshold in each group-round, while 42% fall short, and these rates do not adjust significantly conditional upon success in a previous round. Still, groups do appear to adjust contributions in the direction of increasing payoffs. Total contributions fall an average of 10 units in the subsequent round after exceeding the threshold. Likewise, they tend to rise by 12 units starting from below and 6 units starting at the threshold.

The fact that there is no trend in total contributions or provision rates across rounds, that total contributions are likely to change after a round where the threshold is exactly achieved, that the rate of achieving provision does not change conditional upon success in a previous rounds all lends strong support to the concept of a stochastic steady-state. Furthermore, the tendency of total contributions to decrease when above the threshold in previous rounds and vice versa when they were below is consistent with the motivation of noisy directional learning behind the logit equilibrium (Anderson et al., 2004). Players update choices in the direction of increasing expected payoffs, but they consistently make mistakes with regard to magnitudes. Presumably, this is because they cannot accurately predict the degree to which others will adjust their contributions. The result is a stochastic steady-state in which realized contributions bounce noisily about the

threshold, but a persistent degree of decision error prevents convergence, just as appears in this data.

To summarize, all of the discussion here favors the inclusion of decision errors and altruism in any public goods model. The preference specifications above predict a similar increase in individual contributions relative to self-interested preferences, both below the threshold (away from the inefficient equilibria) and above the threshold ($\alpha = 1, \gamma \sim 1, \beta \rightarrow 0$), so it seems intuitive that incorporating altruism into the QRE model will stochastically increase contributions. This adjustment to incentives would seem a promising route at first blush, given the systematic under-prediction of the QRE model in Tables 4 and 6. A note of caution is in order here, as one should recall that QRE accounts for the feedback loop of stochastic choice on beliefs. While it appears that contributions should be increasing in α and γ , the realization that every other player will increase her contributions may be cause to reduce one's own contributions, especially if $\alpha, \gamma \in (0, 1)$. How these competing incentives net out requires numerically estimating the model itself.

The beauty of the QRE to model is that the stochastic predictions allow for structural estimates of any number of model parameters, in addition to the degree of noisiness, such as the coefficients on altruistic preferences. The danger of QRE is the computation complexity behind estimating the range of fixed-points nec-

essary to build the model and the structural estimators. Anderson et al. (1998) estimate a three parameter model $(\lambda, \alpha, \gamma)$, but the linear utility had a closed-form equilibrium, and the quadratic specification lead to an estimator with a Gaussian distribution. To do a similar analysis with the threshold game would substantially improve this work, but the discontinuous nature of payoffs and the particular design of this experiment (changing states, changing information, changing group-size) make this a computationally heavy lift using currently available tools. For one thing, it is not even clear that pure-altruism or Leontief preferences will satisfy the requirements of a stochastic potential game, while warm-glow can be easily incorporated into utility by $\theta_i(c_i) = v_i + \gamma c_i$. Coincidentally, Offerman and Sonnemans (1996) only found evidence for warm-glow utility in the discrete threshold game, so both the theory and previous empirical work support a two-parameter extension of the model (λ, γ) .

1.5 Conclusion

The purpose of using QRE in this research is straightforward: to obtain theoretical comparative statics that are normally indeterminate, due to the threshold game's multiplicity of equilibria. Using pure strategies under complete information, there are dominated equilibria with very little total contribution and a provision rate of zero, and there are efficient equilibria with total contributions at the

threshold and a provision rate of one. This is true for any group-size, so long as total benefits exceed the threshold. Thus, any predicted difference across group-size using pure-strategy Nash equilibrium requires an ad-hoc switch in equilibrium selection as group-size grows.

Yet, the data presented here show strong evidence against the proposition that players are coordinating on dominated equilibria, as total contributions tend to be near the threshold. In fact, the data suggest that players fail to coordinate, generally, as total contributions admit a stationary degree of noise across all rounds. The logit equilibrium is a stochastic steady state in which players respond to their incentives, albeit imperfectly, so it seems well-equipped to explain coordination failure that tends to be near the threshold. The stationary noise about the threshold in Figures 1.8 and 1.9 appear remarkably consistent with such a state of affairs.

This approach proves fruitful, as the model is able to make the necessary comparative static predictions that are otherwise left ambiguous. QRE accurately predicts the sign of all nineteen of the comparative statics that are significant at the 5% level in this experiment. The model does a reasonable job of capturing the magnitudes of these effects at the group-level; though the same cannot be said about individual contributions. Furthermore, the model tends to systematically under-predict the level of average contributions and over-predict the variance of

contributions, at both the individual and group level. It also tends to under-predict the average provision rate (Figure 1.9); although, it does a good job of explaining over-contributions by low payoff players.

However, these biases in prediction are potentially informative. Golman (2011) proved that the misspecification of homogeneity in the logit parameter introduces downward bias, i.e., estimates will over-state the degree of bounded rationality. In this environment, that is consistent with a lower average and greater variance of contributions for individuals and groups. One obvious solution is to allow the logit parameter to float across payoff types, as this would preserve symmetry. If this approach were successful, it might also be possible to restrict the degrees of freedom across treatments. Thus, heterogeneity in the noise parameter across individuals may improve the accuracy of predictions, while weakening the criticism of over-fitting.

This research has yet to formally incorporate a measure of altruism into the QRE model. This would allow for structural estimation of the altruism parameters. Such motivations are common considerations in any public goods game, and the basic intuition that this will bias contributions upward fits with the bias in levels of current QRE predictions. However, altruism and fairness alone that ignore decision error will be of limited value for this game, as the set of equilibria

Chapter 1. Explaining pledges to threshold public goods with logit equilibrium

tend to include all contributions at or above the threshold; making comparative static predictions impossible.

Chapter 2

Noisy best response and scale

illusion in the market entry game

2.1 Introduction

The market entry game is a simple representation of an important and general payoff structure found in economic life. In this environment, each player faces a binary choice between “entering” or “staying out” of a market. Staying out affords a reservation wage, while entry yields a wage that is decreasing linearly in the total number of entrants. The market capacity is defined as the total level of entry at which payoffs between entering and staying out are equalized for all players. The pure-strategy Nash equilibrium predicts, unsurprisingly, that players will choose to enter the market up to its capacity.

The first and most direct interpretation of this game is the market entry choice faced by firms in microeconomics. In the long run, firms will produce at minimum average cost, suggesting that markets will only have a finite carrying capacity for profitable firms. Likewise, the Cournot model predicts that profits will decrease from monopoly rents to competitive profits as the number of firms grows and market power vanishes. Since each entering firm imposes a negative externality on the payoffs of others, a second interpretation is that of an n -player, binary choice tragedy of the commons game. Lastly, the notion of a congestable public good leads naturally to the classic highway congestion problem with linear time costs. Here, travelers choose between a riskless but circuitous path, such as a beltway, and a risky but direct path, such as a bridge. On the risky route, travel times are increasing in the number of travelers. Nash equilibrium implies that all of the time-saving benefits of the highway will be dissipated by congestion in equilibrium. This is known as the Pigou-Knight-Downs paradox (Arnott and Small, 1994).

However, understanding of observed behavior in this environment is largely incomplete. In the case of homogeneous preferences, the Nash equilibrium concept is silent about exactly which subset of players will enter the market. In fact, any combination of entrance/abstention is acceptable, so long as the total number of entrants is equal to the market capacity. It is reasonable to suspect that

players will fail to coordinate exactly at the market capacity in the absence of communication, especially in a one-shot environment. To complicate matters, there are several mixed-strategy Nash equilibria as well. There is a unique, symmetric mixed-strategy equilibrium in which every player chooses to enter the market with the same probability. There are also a host of asymmetric equilibria in which some players enter or stay out of the market with certainty, while others mix between entry and abstaining with some homogeneous probability.

The selection of an equilibrium that either (a) best represents observed behavior or (b) is socially desirable is far from obvious. Payoff-dominance would suggest that firms with the highest amount of strategic risk should enter the market, while risk-dominance would suggest the opposite. On the other hand, payoff dominance would suggest that travelers with the highest value of time should be sorted onto the bridge, which does make a good deal of sense from a policy perspective. The equilibrium refinement that is desirable or most applicable depends upon the interpretation of the environment.

Given the theoretical indeterminacy of predictions, the Nash equilibrium concept has limited use in predicting behavior here. Thankfully, experimental methods have proven to be an important tool for learning about decision-making in such cases. Due to its wide applicability, this particular game has been studied extensively in the experimental literature (Rapoport, 1995; Rapoport et al., 1998;

Selten et al., 2007), including analysis using models of learning (Rapoport et al., 2002; Erev and Rapoport, 1998; Duffy and Hopkins, 2005), or measurements of overconfidence (Camerer and Lovo, 1999). Most of these analyses focus on perfect versus imperfect information regarding the actions of others, while Camerer and Lovo (1999) focuses on incomplete information regarding one's own type.

This paper focuses on payoff heterogeneity, payoff uncertainty, and group-size effects. Specifically, I examine payoff heterogeneity that takes the form of affine transformations, implying homogeneous preferences for entry. Payoff types are distributed as high or low with equal probability, and each player always knows her own payoff-type. All players know the full state of nature under complete information, while they only know the binomial distribution over others' types under private information. I also vary group-size and the market capacity such that the symmetric, mixed-strategy equilibrium is fixed. These treatments are designed so that there should be no systematic differences across payoff-type, group-size, information completeness, or state of the world so long as individuals behave in accordance with strict best-response.

Given the difficulty of the coordination problem presented to players, there is little reason to believe that perfect coordination upon a single, pure-strategy Nash equilibrium will ever obtain. Previous research has found that decision-makers are noisy and that coordination is unstable in this environment (Hartman,

2007). However, a lack of perfect coordination by no means disproves that players behave in accordance with best-response, given the large number of mixed-strategy equilibria. Under complete information, these mixed equilibria should still exhibit uniform entry rates across payoff-type. Under private information, there is some allowance for heterogeneity in individual entry rates. However, such heterogeneity must leave the best response of all other players unperturbed.

The experimental findings presented here strongly reject these neutral comparative statics, as well as the restrictions that the mixed Nash equilibria impose upon the data. Therefore, an alternative model of noisy decision-making is required. I apply quantal response equilibrium (QRE), which relaxes the assumption of pure best-response and side-steps the issue of indeterminacy. This approach is fruitful in that it provides unique predictions that account for the dozens of systematic differences across payoff-type, information completeness, group-size, and state of the world observed in the experiment. Thus, I find that the noisy behavior observed in these games is both tractable and rationalizable with a model of noisy best-response, while the same cannot be said of pure best-response without the addition of unobserved preference heterogeneity.

It should be no surprise that a model of quantal response behavior matches the data well, so it is important to emphasize what distinguishes this theoretical model from more typical econometric tools. The most comparable technique for

this environment would be a binary response model such as a logit regression. Such models are often applied atheoretically; although, McFadden (1974) showed how they can be motivated through an underlying model of random utility, i.e., unobservable preference shocks. This makes sense for decisions made in markets, where all interaction across individuals is indirectly channeled through the price mechanism. Thus, it is perfect reasonable to assume a unidirectional flow of causality from expected utility to choice probabilities. By contrast, games specify environments where agents interact directly, i.e., every player's expected utility is defined as an explicit function of every other player's choice probabilities. Hence, causality flows in both directions, implying that expected utility and choice probabilities are simultaneously determined. One must apply an equilibrium model to account for this reverberation of noise across individuals; or else risk obtaining biased estimates and predictions.

The remainder of the paper is organized as follows. Section 2.2 presents the theoretical model of the market entry game; including some of the predictions and characteristics of NE and QRE, under both complete and private information. Section 2.3 describes both experimental designs, presents the data, and conducts a battery of tests to demonstrate the relative value of QRE. Section 2.4 discusses the econometric advantage of QRE over the common conditional logit regression. Section 2.5 concludes with remarks for future research.

2.2 Equilibria in the Market Entry Game

The model is defined as follows. There are n players indexed by $i \in \{1, \dots, n\}$. Each player must choose either to “enter” ($s_i = 1$) or “stay out” ($s_i = 0$) of the market. Payoffs to entry are defined linearly, in keeping with the literature:

$$u_i(s_i, s_{-i}) = \begin{cases} a + v_i(c - S) & \text{if } s_i = 1 \\ r & \text{if } s_i = 0 \end{cases} \quad (2.9)$$

where $S = \sum_{i=1}^n s_i$ is the total number of entrants, a is the fixed component of the wage to entry, r is the reservation wage, c is the market capacity, and v_i is player i 's payoff type. In this paper, both experiments restrict $a = r$, implying homogeneous costs to entry.²⁵ The market capacity is defined irrespective of the composition of payoffs.

2.2.1 Nash and Bayes-Nash Equilibria

In this section, I will review the set of Nash equilibria originally discussed in Rapoport et al. (1998), and I will extend these results to Bayesian game with private information. The set of Bayes-Nash equilibria under complete information

²⁵One exception to this rule is the toll treatment in Experiment I that reduces $a < r$. In this case, payoff types (v_i) are homogeneous, so the result is isomorphic to a simple reduction in the capacity, c . Without this restriction, variability in v_i could generate variation in preferences for entry across payoff types.

corresponds to the set of cursed equilibria (Eyster and Rabin, 2005) since each player's payoffs are fully independent of other players' types. This last statement implies that the set of Nash equilibria under complete information are a nested subset of the Bayes-Nash equilibria under incomplete information. Moreover, the set of cursed equilibria implies a testable restriction on the form of individual heterogeneity in entry rates that will be utilized in the results section.

One set of pure-strategy Nash equilibria in this context are rather obvious: namely, any situation such that the market is just at capacity ($S = c$), in which everyone who enters earns at least the reservation wage, while those who stay out strictly prefer to remain out. Additionally, any combination of entry that is one firm short of capacity ($S = c - 1$) is also a Nash equilibrium. This is because all of the firms who choose to enter earn strictly more than the reservation wage, while all of the firms who stay out are indifferent about entering. All of these equilibria are weakly stable under best-reply dynamics.

Solving for the mixed-strategy equilibria is a bit more complicated, but the intuition is straightforward. A player who strictly best responds will be willing to mix between entering and staying out of the market if and only if she is indifferent between these two choices. Therefore, any individual who chooses a mixed strategy must do so because the expected payoff to entry is equal to the reservation wage.

Otherwise stated, a player is only willing to mix if the expected gain to entry is zero.

Although, the solution is identical to the results presented in Rapoport et al. (1998), the set of mixed-strategy equilibria is restated below for the purpose of exposition (see Theorem 1 in appendix A). Assuming that n_{in} and n_{out} players deterministically choose to enter or stay out of the market, respectively, then all remaining players will choose to enter the market with probability:

$$p = \frac{c - n_{in} - 1}{n - n_{in} - n_{out} - 1} \quad (2.10)$$

Importantly, this probability is the same for every player who mixes, regardless of payoff-type. In the special case that all players choose to mix, we have the symmetric equilibrium where everyone chooses to enter the market with probability $p = (c - 1)/(n - 1)$. As this is the only equilibrium in which everyone pursues the same strategy, it represents a focal point.

Given the apparent irrelevance of payoff heterogeneity under complete information, it is natural to wonder what would happen if players were not privy to such information. To be specific, I will continue to assume that each player has private information regarding her own type, v_i . However, I will answer the ques-

tion: “what happens if players only know the ex-ante distribution of other players’ payoff types?”

Each v_i is drawn i.i.d. from a finite set of values $V = \{v_1, \dots, v_K\}$, with a generic probability measure $P(v_k) = q^k$, and I will assume that all of this is common knowledge. Under these conditions, the market entry game is transformed into a Bayesian game, in which the appropriate solution concept for prediction is the Bayes-Nash equilibrium. Therefore, an equilibrium must specify the action that each player will take contingent upon every payoff-type.

All of the Nash equilibria under complete information remain Bayes-Nash equilibria under this definition of incomplete information. To understand why, note that the composition of payoff types was irrelevant under complete information. Therefore, it is simply a matter of transcribing these equilibria into the proper Bayesian notation. For example, the symmetric mixed-strategy equilibrium is defined by the following probability of entry for every player at every payoff type:

$$p_i^k = \frac{c-1}{n-1} \quad \forall i \in I, v_k \in V \quad (2.11)$$

In a Bayesian game, a player is willing to mix between entry and exit if the expected payoff to entry is equal to the reservation wage, where expectations are taken over the type-contingent choice probabilities of others. Therefore, any player

may choose to enter with different type-contingent probabilities, so long as this deviation leaves the expected payoffs of every other player undisturbed relative to the original Nash equilibria with complete information. Thus, the set of Bayes-Nash equilibria extends the set of all Nash equilibria to include any individual heterogeneity that is robust to cursed beliefs (Eyster and Rabin, 2005). To see this, please see Theorem 2 in appendix B; although, the matter is trivial once one recognizes the fact that no individual's payoffs depend upon other players' types within this game. The implication for the binary type space used for experiment II is that any individual may choose to increase her rate of entry as a high (low) type, so long as she reduces her rate of entry as a low (high) type to balance out everyone else's expected payoffs. Importantly, this implies a simple set of testable restrictions on the data that are examined in the last part of section 2.3.

2.2.2 Quantal Response Equilibria

Under the assumption of pure best-response, and without additional restrictions, individual choices remain indeterminate within this environment. While many sensible restrictions could resolve this strategic uncertainty, the data suggest that subjects struggle to anticipate each others' actions correctly. Yet, it is often argued that individuals must somehow form beliefs in order to make decisions (Morris and Shin, 2001; van Huyck and Viriyavipart, 2013). Therefore,

this indeterminacy should be interpreted as a failing of the model, not individual behavior, and another model should be used which provides more meaningful predictions. To that end, I propose using logit quantal response equilibrium, which relaxes the assumption of individual best-response, while maintaining the equilibrium assumption that choices and beliefs be mutually consistent. This approach yields meaningful theoretical prediction that resolve the indeterminacy and strategic uncertainty found under strict best response.

Formally, a logit quantal response equilibrium is defined by the following:

- i. A set of quantal response functions $\{\sigma_i\}_{i=1}^n$ that monotonically transform expected payoffs into choice probabilities according to the logit specification:

$$\sigma_i(s_i; \lambda) = \frac{\exp(\lambda \pi_i(s_i))}{\sum_{s_i} \exp(\lambda \pi_i(s_i))} \quad (2.12)$$

- ii. Expected payoffs that correctly account for the noisy behavior of others:

$$\pi_i(c_i) = u_i(c_i, \sigma_{-i}) \quad (2.13)$$

The exact cause of this deviation from pure best-response may remain unspecified. It may include many mutually consistent alternatives such as bounded rationality (Chen et al., 1997), noisy learning (Anderson et al., 2004), random utility, a

preference for unpredictability or a cost to accuracy, fairness or altruism. The key feature that distinguishes logit QRE from, say, a logit regression model is the requirement that beliefs correctly account for the effect that each individual's noisiness has on every other player. Therefore, expected payoffs and choice probabilities are simultaneously determined, and the QRE explicitly accounts for the reverberation of noise across individuals in addition to individual noisiness, itself.

In games with a unique Nash equilibrium, the logit QRE is unique for every value of $\lambda \in [0, \infty)$. This is not necessarily the case in games of multiple equilibria, such as the market entry game. However, it is clear that QRE is always unique for $\lambda = 0$, where every player must be randomizing uniformly over actions. In fact, McKelvey and Palfrey (1995) demonstrate that this uniqueness holds for every $\lambda \in [0, \lambda^*]$, where λ^* is some constant that is specific to the game, as well as the particular parameterization. Moreover, the authors prove that there exists a continuous correspondence of equilibria connecting uniform mixing at $\lambda = 0$ to a generically unique Nash equilibrium as $\lambda \rightarrow \infty$. In the event that a game has multiple equilibria and is perfectly symmetric, as will be the case in experiment I and certain instances in experiment II, this uniqueness may no longer hold. Fortunately, this issue does not apply to the market entry game, as will be explained in the next section where I describe the limiting equilibria. Therefore,

the principal branch is strictly unique, and provides meaningful predictions for my purposes.²⁶

While the logit specification prevents closed-form solutions to this fixed point, modern homotopy methods may be implemented numerically to solve for the principal branch of the logit equilibrium correspondence. I use the algorithm described in Turocy (2005) to numerically calculate the principal branch of logit QRE under complete information. Using the celebrated Harsanyi transformation to conceptualize the game of incomplete information as a game of imperfect information regarding “Nature”, I am able to use the extensive-form analogue to QRE known as “Agent QRE” (McKelvey and Palfrey, 1998) to guarantee existence of a solution. I then solve for these equilibria using the algorithm described in Turocy (2010). Both codes are extensions of the algorithm incorporated into The Gambit Project (McKelvey et al., 2013); a set of open-source tools for computational game theorists. Many thanks to Theodore L. Turocy for generously sharing this code, which is designed to handle games with group-sizes larger than two. The results of these numerical calculations are then used to construct the structural maximum likelihood estimators for the logit parameters. Since the model makes stochastic

²⁶The principal branch is the only subset of quantal response equilibria that is applied throughout the literature. The reasons for this include: (i) it is the only set of equilibria that we know for certain exists, (ii) it is the only set of equilibria that can possibly exist for $\lambda \in [0, \lambda^*]$, and (iii) it is the only set of equilibria that we have a universal method for solving numerically.

predictions, the likelihood function for this estimator is defined implicitly by the quantal response functions as stated below.

$$L(\lambda; s_i) = \frac{1}{N} \sum_{i=1}^N \log \sigma_i(s_i; \lambda) \quad (2.14)$$

2.2.3 Quantal Response Equilibrium Predictions

One key feature of QRE is that choice probabilities are sensitive to payoff magnitudes, sometimes referred to as “scale illusion”. In the pure-strategy equilibrium of the market entry game, where the total number of entrants equals the market capacity, the average gain to entry is fully dissipated. Hence, there should not be any difference in entry rates between high and low payoff players, regardless of scale illusion. However, if such an equilibrium does not attain - not even on average - then either the market will tend to be systematically above or below its capacity. Supposing that the market is systematically under capacity, there would be a consistent and positive average gain to entry, and this gain would be larger in magnitude for high payoff players than it would be for low players. We might intuitively expect to find that high payoff players enter more often, since they have more to gain. The QRE does in fact predict under entry relative to capacity

in the absence of a toll, given the parameterizations used here. Therefore, the QRE also predicts a higher rate of entry for high payoff players.²⁷

Armed with this simple intuition, and the knowledge that QRE is symmetric by payoff-type, it is easy to work through the logic of the limiting logit equilibria as well. In the case of homogeneous payoffs, it is clear that every player must enter with identical probabilities. The only symmetric Nash equilibrium of the game is the mixed-strategy where each player chooses enter with probability $p = (c - 1)/(n - 1)$, so this is the limiting equilibrium selected in experiment I as well as the symmetric states of nature in experiment II.²⁸ A similar logic applies to the other states of nature as well. For example, suppose the game has four players with a market capacity of three, as will be the case in the small group treatment. If the realized state has three high players and one low player, then the limit will be the pure-strategy equilibrium where the three high types enter and the low type stays out. By contrast, suppose there is one high type and three low types. In that case, the selected equilibrium will have the high type player entering, while the remaining low types enter with probability $p = (c - 2)/(n - 2)$; corresponding to one of the asymmetric mixed-strategy equilibria. The pattern is quite clear:

²⁷Experiment I makes use of a toll, but given the homogeneous payoffs, I cannot test the reverse logic regarding payoff types. However, the QRE predicts over entry in this case. As will be shown in the next section, this is exactly what I find in the data.

²⁸A word of caution here. In a perfectly symmetric game with multiple equilibria, even the principal branch may be non-unique for some $\lambda^* > 0$, which would imply multiple limiting equilibria. The reason this is not the case here is because there is one-and-only-one symmetric Nash equilibrium. That is a feature specific to this particular environment.

the Nash equilibrium selected by the limit of QRE will be one in which as many high payoff players enter as possible, while still respecting symmetry.

Of course, not every prediction made by the QRE is different from the standard Nash predictions. Experiment II is designed so that any pure information and group-size effects should be negligible according to pure best-response. Quantal response behavior does not have anything substantively different to say about these variables, individually. Still, the QRE predicts that the difference in entry rates between high types and low types contracts as group-size grows. So while both models do not predict a group-size effect, per se, quantal response behavior does lead to an interaction effect between group-size and scale illusion that is absent under best response. As will be discussed in section 2.4, this effect is also missed by logit regressions that use the average payoffs experience by subjects within the experiment. The key difference between these approaches is that a regression assumes unidirectional causality from expected payoffs to choice probabilities, while logit QRE accounts for the fact that causality flows in both directions simultaneously. Hence, the QRE is doing more than simply fitting an econometric model and re-stating the observation that scale illusion matters.

2.3 The Experiments

This section details the design and results of two experiments. The first experiment uses homogeneous payoffs to create a baseline measure of behavior, and tests the effect of introducing a toll, i.e., reducing the fixed wage to entry from $a = r = 0$ to $a < 0$. The second experiment uses heterogeneous payoff types, but with homogeneous preferences, by maintaining $a = r = 0$ while allowing each player v_i to vary. This experiment tests for scale illusion, payoff uncertainty, and group-size effects. It is necessary to run a new set of experiments to make sure that all comparisons across treatments are comparable, since the QRE model is sensitive to affine transformations of payoffs. Both experiments were conducted in the Experimental and Behavioral Economics Laboratory (EBEL) at UC Santa Barbara using the z-Tree software package (Fischbacher, 2007). Subjects were randomly recruited from a subset of the undergraduate population that has registered their desire to participate in experiments.

The first experiment was originally presented in Hartman (2007). It was designed to answer questions about traffic networks and congestion externalities. However, the market entry game is isomorphic to the Pigou-Knight-Downs paradox (Arnott and Small, 1994). This experiment is useful as a baseline scenario where subjects face perfectly symmetric incentives, remain in fixed groups, and

observe the total number of entrants between rounds. The environment encourages learning, which gives participants the best chance of achieving either a pure-strategy equilibrium at the capacity, or the symmetric, mixed-strategy equilibrium. In addition, the experiment makes use of a toll to nudge behavior towards a more socially desirable outcome.

I conducted the second experiment. In this case, the experiment was designed to focus on group-size and information effects. The private information treatment makes use of Bayesian games, so subjects values have to be re-drawn each round. In this experiment, subjects played ten practice rounds where they observed the total number of entrants after each round, and twenty potentially paid rounds of each information treatment. To avoid the repeated-game cooperation, groups are rotated in the practice rounds, and this practice is maintained in the potentially paid rounds to avoid confusion. In this case, the deck is stacked against convergence to a pure-strategy equilibrium.

There is little evidence for convergence to pure-strategies in either experiment. In fact, half of all subjects switch choices in a round following an equilibrium outcome, and this rate increases for the second experiment. These findings are perfectly consistent with Selten et al. (2007), who study the traffic congestion problem and find that choice frequencies are unaffected by imperfect information between rounds, but switching rates increase as information decreases. It seems

that noisiness and coordination failure are inherent features of the data produced using this incentive structure. This fact is not terribly sensitive to experimental design choices regarding perfect/imperfect information. All of this supports the relevance of QRE, which is successful at explaining the toll effects in experiment I as well as the interaction of group-size and payoff-type effects in experiment II.

2.3.1 Experiment I

The baseline experiment includes ten sessions. Session nine had only seventeen subjects, while every other session included eighteen subjects. In all cases, the market capacity is fixed to eleven. Each session includes two treatments: in the first, market entry is free, while in the second, entry requires paying a toll. The toll is designed to reduce the number of entrants to the socially optimal level of five. Because values are homogeneous in this scenario, adding the toll is isomorphic to simply reducing the market capacity by six. Sessions include twenty paid rounds for each treatment, and final payments are determined as the sum over all paid rounds. Average payments were \$12 to \$15 per subject for sessions that lasted about an hour.

Since payoffs are aggregated across all 20 rounds in each treatment to determine payments, I shall discuss the model's parameters in terms of these totals. In both cases the market capacity is fixed to $c = 11$ and the payoff magnitudes

are $v_i = 4$ for every player i . In the free treatment, the reservation wage and the fixed wage to entry are equal and normalized for simplicity to $a = r = 0$. The toll is then defined as a reduction of the fixed wage to entry to $a = -24$. As stated above, this is equivalent to maintaining $a = r = 0$ while reducing the market capacity to the socially optimal level of $c = 5$.

In the first treatment where entry is free, the set of pure-strategy Nash equilibria includes any combination of eleven entrants. Given a group-size of eighteen, the underlying rate of entry should be 0.611. However, perfect coordination on a pure-strategy Nash equilibrium is especially difficult in this environment, where some players must enter while others must stay out. The only symmetric equilibrium is a mixed-strategy with an underlying entry rate of 0.589. Taking all asymmetric mixed-strategy equilibria into consideration as well, the observed average rate of entry may vary from 0.556 to 0.611. Thus, the introduction of mixed-strategies should, in theory, allow for average entry rates below the pure-strategy level.

The same is true in the second treatment. If the toll reduces the pure-strategy Nash equilibrium number of entrants to five, then one would expect an average entry rate of 0.278. Supposing that the symmetric, mixed-strategy equilibrium is focal, the average entry rate drops to 0.235. The full range of equilibria, including the asymmetric mixed-strategies, yields an average entry rate varying from 0.222

to 0.278. In both cases, the average entry rate under pure strategies represents an upper bound (see Theorem 3 in appendix A).

Table 2.1: Experiment I: Average Entry by Treatment

Group	Free	Toll
1	0.594 (0.116)	0.325** (0.110)
2	0.617 (0.115)	0.311 (0.109)
3	0.550 (0.117)	0.308 (0.109)
4	0.606 (0.115)	0.286 (0.107)
5	0.614 (0.115)	0.317 (0.110)
6	0.628 (0.114)	0.322** (0.110)
7	0.603 (0.115)	0.317 (0.110)
8	0.603 (0.115)	0.333** (0.111)
9	0.635 (0.117)	0.329 (0.114)
10	0.631 (0.114)	0.306 (0.109)
All	0.608 (0.036)	0.315** (0.035)

Table 2.1 presents the estimated average entry rates observed in this experiment. The results are presented for each session, as well as the whole dataset. Column two corresponds to free entry, with a capacity of eleven, while column three represents the toll for entry, in which the capacity is tacitly reduced to five. The null hypothesis tested in this table is that the average entry rate is equal to its upper-bound, pure-strategy Nash equilibrium rate. The estimated entry rates never differ significantly from the Nash prediction in the free treatment. Surpris-


ingly, estimates vary significantly from the pure-strategy Nash prediction in the toll treatment. The pure-strategy prediction is an upper-bound, and introducing mixed-strategies only lowers the expected entry rate.

This could be explained using disequilibrium. For example, suppose each subject believes her group has coordinated upon an equilibrium in which she may enter with certainty. Then the experimenter should expect an entry rate of one, well above any Nash prediction. However, such an explanation is undercut by learning. In this baseline experiment, subjects are in fixed groupings, and they are able to observe the total number of entrants after each round. There is no evidence of such costly errors being sustained across rounds.

However, there is little evidence of convergence either. Almost half of all subjects switch their choice in the subsequent round after an equilibrium has been reached, implying little hope for eventual convergence to a pure-strategy equilibrium. If learning lead to convergence on a Nash equilibrium, the average entry rate should trend across rounds towards the range of rates predicted by the pure- and mixed-strategy equilibria. The data admit a stable and consistent degree of noisy which is centered about the capacity for the free case, but it is biased above the predicted entry rates in the toll case. Figure 2.1 shows that the majority of individuals behaved “as if” they were mixing. There appears to

be significant heterogeneity in entry rates. With only twenty observations per individual, it is unclear how much of the supposed heterogeneity is simply noise.

Figure 2.1: Individual Entry Rates



One natural explanation for the lack of convergence is bounded rationality. Players are able to correctly perceive their incentives, however, they remain uncertain about the actions of others. In a boundedly rational equilibrium, where action uncertainty is left partially unresolved, players would never be able to settle on a pure-strategy equilibrium. The logit form of the QRE is especially well-suited to this task. Starting from $\lambda = 0$, the entry rate will always be 50%, and the limiting logit equilibrium will be the symmetric, mixed-strategy Nash equilibrium. Since the symmetric equilibrium is above 50% in the free case and below 50%

in the toll case, this form of bounded rationality will predict under entry in the former treatment and over-entry in the latter.

Estimating the logit parameter for the entire dataset, with only one degree of freedom across both treatments, I find a sensitivity estimate of $\hat{\lambda} = 0.175$, with a standard error of 0.02. The corresponding entry rate for the free entry treatment is 0.566, which is a little below the estimates found in the experiment. However, it is greater than the estimated rate for session 3, and it is not significantly different from the average rate of 0.608 across all sessions. Likewise, the corresponding entry rate for the toll entry treatment is 0.302. Unsurprisingly, this is not significantly different from the overall rate of 0.315, which was significantly greater than the pure-strategy upper-bound rate for Nash equilibria. Thus, QRE appears to explain over-entry in the presence of a toll. While a small amount of under-entry is potentially consistent with mixed-strategy Nash equilibria, fairness, altruism, or efficiency, over-entry benefits neither the individual nor the group. Intuitively, it is a pure example of a decision error.

2.3.2 Experiment II

The data presented here span eight experimental sessions. Four sessions include four groups of four subjects with a market capacity of three, and four sessions include two groups of seven subjects with a market capacity of five. In

each session, subjects participated in fifty rounds of play: ten practice rounds and two sets of twenty potentially paid rounds; one for each information treatment. Subjects' types are distributed Bernoulli, with each player chosen to have either a high stakes or low stakes with equal probability. These were re-drawn i.i.d. for each round of the experiment. Under complete information, each subject knows whether she is high or low, as well as the total number of high and low players in her group. Under incomplete (private) information, each player knows whether she is high or low, but only knows that everyone is high or low with equal probability.

Table 2.2: Experiment II: Payoffs

Entrants	Small			Large		
	Low	High	Stay Out	Low	High	Stay Out
0	-	-	15	-	-	15
1	25	35	15	25	35	15
2	20	25	15	22.5	30	15
3	15	15	15	20	25	15
4	10	5	-	17.5	20	15
5	-	-	-	15	15	15
6	-	-	-	12.5	10	15
7	-	-	-	10	5	-


The payoff structure to participants in this game is displayed in Table 2.2, and the half of this table that corresponded to any session's group-size was included in the experimental instructions. The values were chosen so that the reservation wage, which is the expected earnings for each subject if they coordinate, is equal to

\$15. Furthermore, I wanted to ensure that the minimum and maximum payments were equal across group-size, for any given payoff type. Lastly, I wanted to ensure that the minimum earnings were equal to \$5, the minimum show-up fee in the experiment. The set of Nash equilibria for this parameterization are presented in Table 2.3. The principal branch of QRE is plotted in Figure 2.2 for the state of nature with two low risk players. Notice that, under complete information, increased responsiveness to payoffs actually increases the noisiness of low risk players by pushing the entry rate towards $1/2$. Thus, λ is a relative measure of payoff responsiveness, not an absolute measure of noise.

Table 2.3: Nash Equilibria

(a) Small			
Nash Eqm	n_{in}	n_{out}	p
Pure-Strategy	3	1	-
	2	2	-
Symmetric Mixed-Strategy	-	-	2/3
Asymmetric Mixed-Strategy	1	0	1/2
(b) Large			
Nash Eqm	n_{in}	n_{out}	p
Pure-Strategy	5	2	-
	4	3	-
Symmetric Mixed-Strategy	-	-	2/3
Asymmetric Mixed-Strategies	3	1	1/2
	3	0	1/3
	2	1	2/3
	2	0	1/2
	1	1	3/4
	1	0	3/5
	0	1	4/5

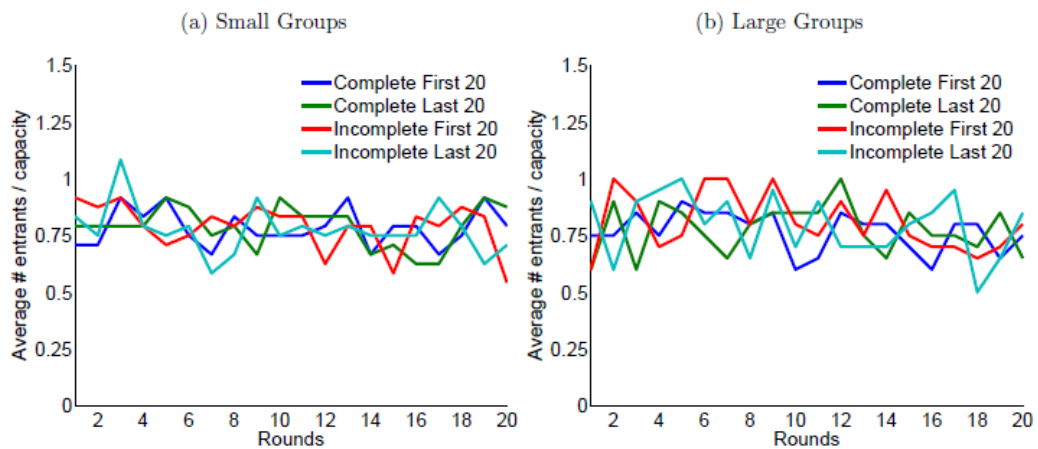
Figure 2.2: QRE: Principal Branch



At the group level, I find that the average entry rate, normalized by the market capacity, is unaffected by group-size, information completeness, or experimental rounds. Among the four sessions for each group-size, two encountered the complete information treatment first, while two faced it second. There do not appear to be any order or session effects. While the plots in Figure 2.3 are consistent

with both noisy best response and mixed-strategy equilibria, they represent strong evidence against pure-strategy equilibria. They unambiguously represent coordination failure in the strict sense.

Figure 2.3: Average Entry Relative to Capacity



Analysis Across States

Averaging across states of nature reveals a more nuanced story, as can be seen in Table 2.4. I observe significant differences across player-type and states of nature under complete information. High stakes players enter more often than low stakes players, and this tendency moves in inverse relation to the total number of high stakes players. Under incomplete information, I merely observe the overall tendency for high stakes players to enter more often than low stakes players. As can be seen in Table 2.5, all of these comparative statics are perfectly consistent

with the predictions of logit QRE at the estimated value of the logit parameter, regardless of whether this parameter is allowed to vary across experimental treatments. This is because the choice probability to enter must be proportional to the expected payoff of entry, which is greater for high risk players and may only be conditioned on the full state of the world under complete information.

Table 2.4: Observed Entry Rates Across States

State Realized	Small Complete Low	Small Complete High	Small Incomplete Low	Small Incomplete High	Large Complete Low	Large Complete High	Large Incomplete Low	Large Incomplete High
0 low types	-	0.550 (0.056)	-	0.652 (0.05)	-	0.381 (0.109)	-	0.571 (0.202)
1 low types	0.567 (0.053)	0.511 (0.03)	0.541 (0.054)	0.592 (0.031)	0.889 (0.076)	0.519 (0.048)	0.667 (0.167)	0.648 (0.066)
2 low types	0.548 (0.033)	0.752 (0.029)	0.570 (0.032)	0.620 (0.031)	0.565 (0.074)	0.574 (0.046)	0.579 (0.057)	0.511 (0.036)
3 low types	0.548 (0.033)	0.763 (0.049)	0.559 (0.033)	0.635 (0.056)	0.58 (0.04)	0.505 (0.035)	0.601 (0.042)	0.587 (0.036)
4 low types	0.526 (0.058)	-	0.529 (0.061)	-	0.471 (0.042)	0.724 (0.044)	0.538 (0.044)	0.535 (0.05)
5 low types	-	-	-	-	0.550 (0.046)	0.646 (0.070)	0.592 (0.045)	0.688 (0.068)
6 low types	-	-	-	-	0.433 (0.092)	0.800 (0.200)	0.619 (0.076)	0.571 (0.202)
7 low types	-	-	-	-	0.143 (0.097)	-	0.571 (0.137)	-

Table 2.5: Predicted Entry Rates Across States

(a) One parameter total

State Realized	Small Complete Low	Small Complete High	Small Incomplete Low	Small Incomplete High	Large Complete Low	Large Complete High	Large Incomplete Low	Large Incomplete High
0 low types	-	0.596	-	0.610	-	0.577	-	0.610
1 low types	0.539	0.607	0.556	0.610	0.537	0.579	0.556	0.610
2 low types	0.546	0.626	0.556	0.610	0.538	0.607	0.556	0.610
3 low types	0.555	0.647	0.556	0.610	0.550	0.613	0.556	0.610
4 low types	0.566	-	0.556	-	0.553	0.619	0.556	0.610
5 low types	-	-	-	-	0.558	0.629	0.556	0.610
6 low types	-	-	-	-	0.547	0.600	0.556	0.610
7 low types	-	-	-	-	0.552	-	0.556	-

(b) One parameter per treatment

State Realized	Small Complete Low	Small Complete High	Small Incomplete Low	Small Incomplete High	Large Complete Low	Large Complete High	Large Incomplete Low	Large Incomplete High
0 low types	-	0.607	-	0.614	-	0.577	-	0.584
1 low types	0.537	0.628	0.558	0.614	0.537	0.579	0.542	0.584
2 low types	0.546	0.653	0.558	0.614	0.538	0.582	0.542	0.584
3 low types	0.559	0.692	0.558	0.614	0.540	0.586	0.542	0.584
4 low types	0.579	-	0.558	-	0.542	0.591	0.542	0.584
5 low types	-	-	-	-	0.545	0.595	0.542	0.584
6 low types	-	-	-	-	0.547	0.600	0.542	0.584
7 low types	-	-	-	-	0.550	-	0.542	-

Analysis of Heterogeneity under Private Information

When information is incomplete, i.e. players only observe private information regarding own payoff-type, but the ex-ante distribution of types is common knowledge, the Bayes-Nash equilibrium concept allows for heterogeneity in entry rates across any one player's payoff type, which further implies heterogeneity across individuals generally. However, this divergence must be robust to cursed beliefs, which imposes a testable restriction upon the data.

The constraint concerns the relative entry rates of each individual across payoff types. According to the parameterization used here, a mixed-strategy must satisfy the following condition in order to be consistent with Bayes-Nash equilibrium (see Theorem 4 in appendix A):

$$p_l^i = \alpha + \beta p_h^i \quad \text{where} \quad \alpha = \frac{1}{q} \cdot \frac{c - n_{in} - 1}{n - n_{in} - n_{out} - 1} \quad \text{and} \quad \beta = -\frac{1 - q}{q} \quad (2.15)$$

For a group-size of four, and equal probability of types, the parameters are: $\alpha \in \{1, 4/3\}$ and $\beta = -1$. These are testable restrictions on the data. I estimate a number of restrictions on the full set of observations, as well as a restricted sample that excludes those who entered or stayed out with certainty. Results are reported in the Table 2.6.

Table 2.6: Test of BNE restriction

(a) Small

Restriction	Sample Size	$\hat{\alpha}$	$\hat{\beta}$
None	64	0.538 (0.102)	0.045 (0.146)
$\alpha = 1$	64	-	-0.550 (0.074)
$\alpha = 4/3$	64	-	-1.000 (0.090)
Restricted Data	55	0.714 (0.099)	-0.343 (0.154)
Restricted Data & $\alpha = 1$	55	-	-0.741 (0.071)
Restricted Data & $\alpha = 4/3$	55	-	-1.204 (0.087)

With only one exception, the null hypothesis that $\beta = -1$ is rejected at the 5% level. The one exception is for the full data, with $\alpha = 4/3$. In that particular case, the estimate is spot on. Thus, I do not definitively reject the BNE mixed-strategies. Nonetheless, in both specifications where α is a free parameter, I strongly reject both that $\alpha = 1$ or $\alpha = 4/3$ and $\beta = -1$. Therefore, these estimates present reasonable evidence that the mixed-strategies seen in the data are not consistent with BNE. Therefore, players' mixing rates are not consistent with indifference or best-response behavior. The entry rates observed here give rise to positive expected payoffs from entry; an opportunity that appears to be under-exploited.

Analysis of Payoffs

This deviation from pure best response appears to be ex-ante welfare-improving. Under pure strategy equilibria, everyone earns the reservation wage. In all of the mixed-strategy equilibria, players are willing to mix because the ex-ante expected payoff from entering is equal to the reservation wage; inducing indifference. Thus, the only players who are predicted to earn more than the reservation wage are those who enter with certainty, in an asymmetric mixed-strategy equilibrium. By contrast, QRE allows for the ex-ante expected payoff from entry to depend upon payoff-type, state of nature, and the degree of noise. Table 2.7 reports the observed average payoffs in excess of the reservation wage. Likewise, I report the range of expected payoffs predicted by the asymmetric, mixed-strategy equilibria and the QRE.

Table 2.7: Average Observed Benefit to Entry

Group-size	Information	Payoff-type	All Data	Restricted	Mixed NE	QRE
Small	Complete	Low	0.46*** (0.12)	0.30** (0.12)	0.63	0.32-0.74
Small	Complete	High	2.07*** (0.26)	1.90*** (0.28)	1.25	1.06-2.24
Small	Incomplete	Low	0.65*** (0.13)	0.50*** (0.13)	0.63	0.67
Small	Incomplete	High	1.39*** (0.27)	1.06*** (0.27)	1.25	1.48
Large	Complete	Low	0.90*** (0.10)	0.94*** (0.11)	0.09-0.71	0.72-1.00
Large	Complete	High	2.03*** (0.20)	2.09*** (0.22)	0.17-1.43	1.63-2.21
Large	Incomplete	Low	0.72*** (0.11)	0.65*** (0.11)	0.09-0.71	0.84
Large	Incomplete	High	1.43*** (0.22)	1.29*** (0.23)	0.17-1.43	1.82

Earnings are statistically significantly higher than the reservation wage across all treatments. Yet, most subjects did not enter nor stay out with certainty. This implies either (i) the absence of pure best-response behavior, (ii) players are best-responding to mistaken beliefs, or (iii) some other criteria affects decision-making beyond pure self-interest. The QRE predicts the average payoffs quite well under complete information, as well as under incomplete information for small groups. Meanwhile, the mixed-strategy equilibria predict better for large groups with incomplete information.

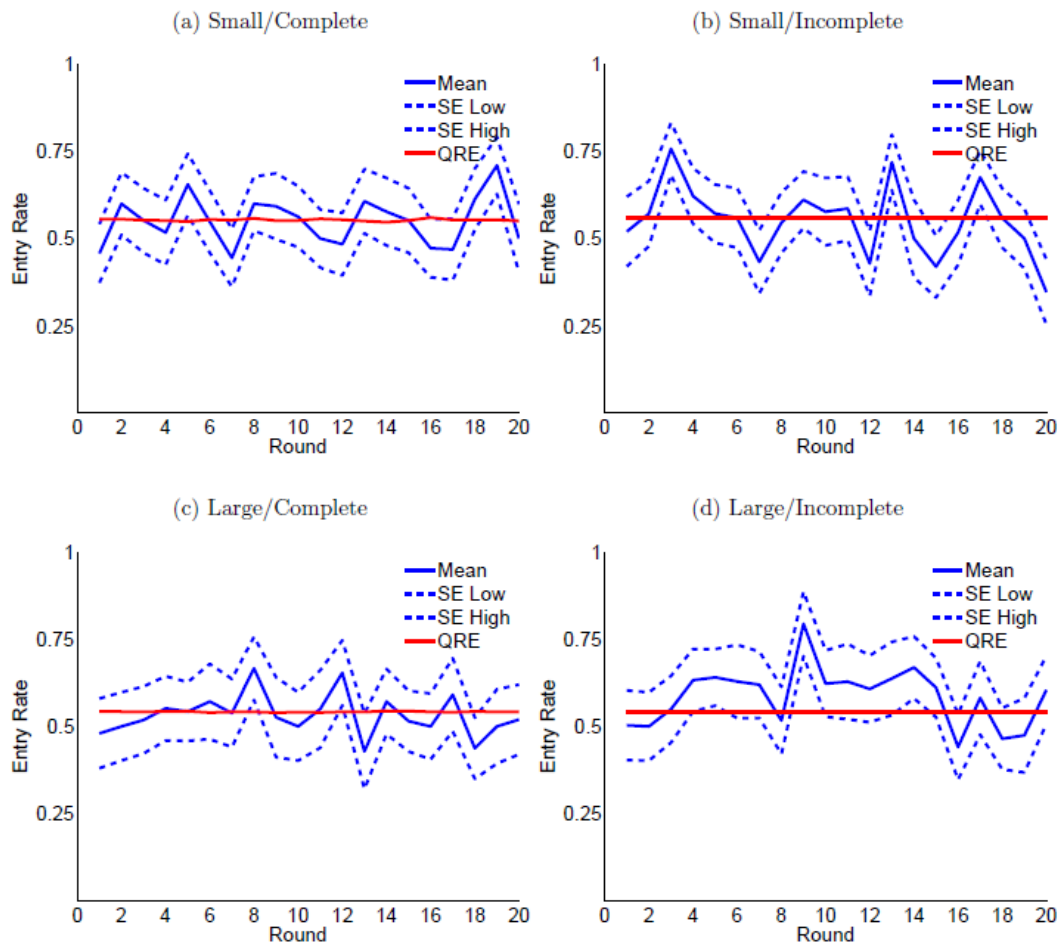
Analysis Across Rounds

I previously decomposed the rates of entry across states of nature and payoff types. Sample means were estimated across subjects as well as rounds. That approach leads to multiple observations per subject because states and types are often repeated across rounds. For example, a subject may find herself to be one of three low payoff players in round 1 as well as in round 5 of a session. By contrast, each subject may only make one choice per round, so a decomposition across experimental rounds will only contain one choice per subject.

The plots in Figures 2.4 and 2.5 admit no trend in entry rates across rounds, regardless of payoff-type, group-size, or information treatment. Thus, it would appear that players do not improve at coordination over time. This is consistent with the design of the experiment. Group matching is rotated each round, and

outcomes are not revealed, so as to suppress the tacit coordination that might be expected in a repeated game environment. To the extent that QRE is useful as a descriptive model of during coordination failure, this is a desirable outcome.

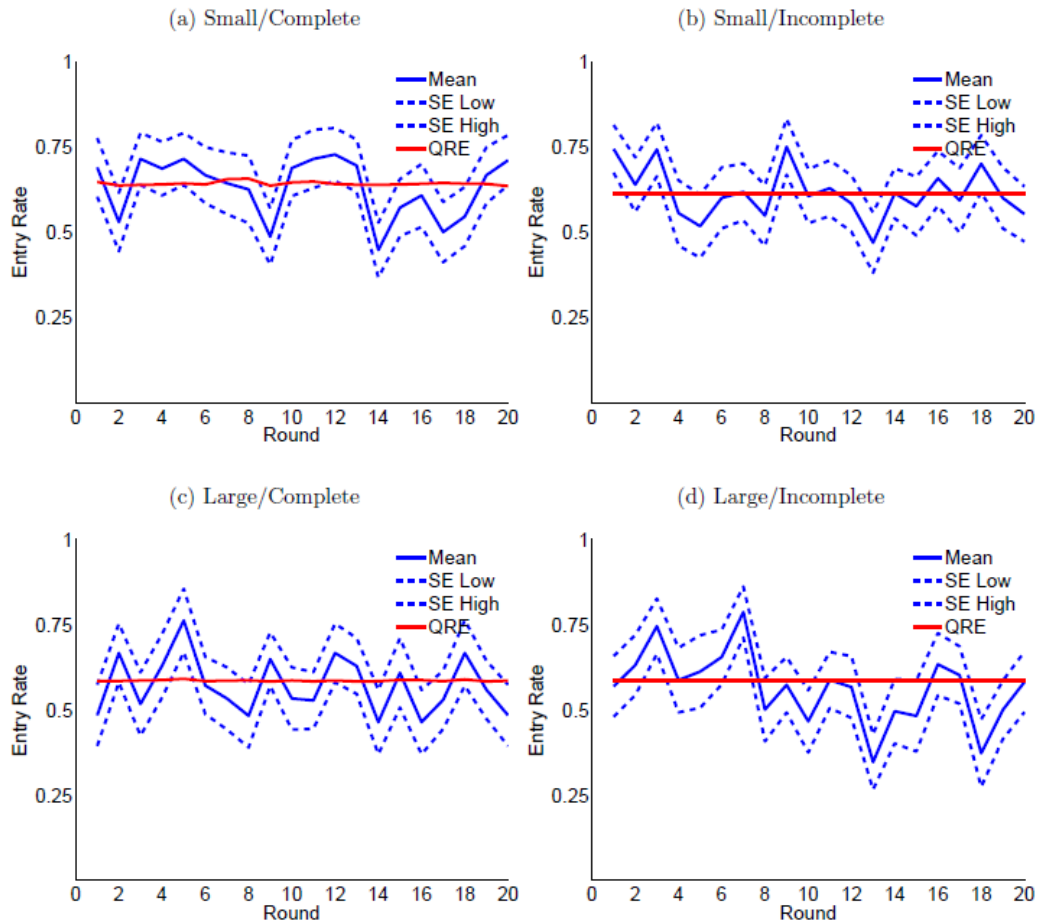
Figure 2.4: Low Entry Rates



QRE predictions fall within one standard error of the estimated entry rates with only rare exception. Furthermore, there is little evidence that the random

draws of nature have a meaningful impact on overall predictions or outcomes. Variation in states has only a slight effect on QRE predictions under complete information, and, of course, no effect under incomplete information. Therefore, the variation in estimated entry rates from round-to-round reflects statistical noise, as there are only 40 subjects per treatment-round. This persistent noise is indicative of the struggle to solve the underlying coordination problem.

Figure 2.5: High Entry Rates



One concern with this data is whether or not players' choices are independent across rounds. Both QRE and BNE mixed-strategies assume that draws are independent. The experiment is designed so that each round will function as a single shot game. However, given the difficulty of the coordination problem, an individual may play a heuristic that exhibits persistence across rounds. I exam-

ine the average entry rates per round, in order to find any persistence that may exist at the group level. So long as the group as a whole is making independent draws, then either model presents a reasonable “as if” description of behavior at the group level.

Table 2.8: Tests for Persistence

Group-size	Information	Auto-Low	Auto-High	Cross
Small	Complete	-0.081 (0.233)	0.173 (0.242)	-0.093 (0.188)
Small	Incomplete	0.040 (0.275)	-0.047 (0.226)	0.181 (0.314)
Large	Complete	-0.243 (0.230)	-0.185 (0.238)	-0.094 (0.166)
Large	Incomplete	0.185 (0.232)	0.207 (0.237)	0.034 (0.185)

I estimate an AR(1) model²⁹ for the entry rates, averaged within a round, for each payoff type in each treatment. It is clear from the results displayed in Table 2.8 that there is no persistence in entry rates at the group level. Likewise, there is no cross-type correlation across rounds, i.e., if high payoff players enter at a greater than average rate, this has no impact on the entry rates of low payoff players. Thus, at the group level, individuals are not responding to their own previous choices, nor are they responding the changes in the choices of others. This suggests that models in which players randomize with independent draws will be suitable approximations of the behavior observed in the laboratory.

²⁹These regressions included a constant to absorb the positive mean in the data. I have examined these data with an AR(2) model as well, and found no persistence in either case.

Analysis Across Individuals

Upon close inspection of the data, there is an interaction effect between payoff type and group-size. For small groups with complete information, high types enter 63% of rounds, while low types enter in 55% of rounds. However, with incomplete information (private knowledge regarding one's own type, common knowledge regarding the ex-ante distribution of others' types), high types enter 62% of the time, while low types enter 56% of the time. In both cases, these differences are significant at the 5% level. It bears repeating that the set of equilibria includes asymmetric mixing, with some players choosing to enter or stay out with certainty. For robustness, I collect a restricted estimate that excludes those subjects whose entry rate was either 0% or 100%. Under this restriction, high types and low types enter at a rate of 62% and 48%, respectively, for complete information and 58% and 49%, respectively, for incomplete information. Thus, the restriction only strengthens the argument that payoff types matter for small groups.

By contrast, differences in entry rates disappear in both information treatments with large groups, regardless of restrictions on the data. These results are loosely consistent with QRE, which predicts a narrowing in differential rates of entry among large groups. However, failing to reject a hypothesis is not the same as accepting the alternative. Therefore, it is unclear whether there is no difference


in entry rates with larger groups, or whether the sample size is simply too small to effectively measure the smaller differences that may exist.

Looking at the individual level, however, I find a high degree of heterogeneity in entry rates. The dispersion in entry across individuals is wider than anticipated by either theory. It is not clear whether this variance is a rejection of both theories or symptom of small sample size. Since each player's type is re-drawn each round, some subjects were only able to make a small number of decisions as a given type. For example, some participants only made four decisions as a low type. In that case, our estimate of that subject's entry rate as a low type is predominantly statistical noise. This may be why I observe entry rates of almost every decile (10%, 20%, 30%, ...), for all player types and all treatments.

In support of the Nash concept, there also appear to be a number of subjects who always entered or stayed out. To illustrate this fact, note that more than 20% of subjects always entered when low or high in the small-group, incomplete-information treatment. However, only half of these subjects (roughly 10%) always entered as both types. Given the entry rates predicted by QRE, we would reject such an event, even at the 1% level. Yet, the only asymmetric mixed strategy with small groups would require one out of four players to enter with certainty, which is well above the level observed here. As is usually the case in experiments, there remains much room for debate regarding individual behavior. It is possible that

a subset of subjects were best responding to the belief that they had coordinated on an asymmetric, mixed-strategy equilibrium in which they should enter or stay out with certainty.

Figure 2.6: Average Entry Rates



It is possible to cut through much of this individual heterogeneity using proper regression techniques. I estimate³⁰ a fixed-effects linear probability model of entry:


$$P_i(s_i = 1) = \alpha_i + \beta \cdot High_i + \varepsilon_i \quad (2.16)$$

The additional probability of entry by high payoff players is represented by β , and I estimated a value for each treatment. In the simple model, I restricted $\alpha_i = \alpha \forall i$. This is equivalent to the differences show in the bar graphs of Figure 2.6. I also estimate the fixed-effects model, allowing α_i to pool by subject or round. Pooling by subjects allows the model to absorb the large heterogeneity in entry rates observed in Figure 2.7, while pooling across rounds controls for the significant variation in entry across rounds found in Figures 3 and 4.

The results for the simple model, as well as both fixed-effects models are highly robust. They strongly suggest that high types enter more often than low types, when groups are small. However, this difference vanishes for large groups. This is consistent with the entry differentials predicted by QRE. It is not obvious what other models could re-produce this comparative static of coordination failure.

³⁰These results are for the full data set. I have estimated these results, restricting attention away from those who always enter or stay out, and have found similar results.

Figure 2.7: Individual Entry Rates



2.4 Discussion

One central argument advanced in this research is that the rate of entry by individuals will differ across payoff type. This stands in contrast to all of the Nash equilibria under complete information, which predict no difference in entry rates

across those individuals who are mixing. Likewise, Bayes-Nash equilibria allow for different entry rates under incomplete information, but these equilibria imply restrictions that are strongly rejected by the data. Subjects are not expected to solve the equilibrium selection problem, so these discrepancies are presumably characteristics of coordination failure. QRE predicts coordination failure, since players are assumed to be noisy in their decision-making. Likewise, it predicts differential rates of entry, since the probability each strategy is chosen is in some proportion to the expected payoff. Therefore, this model is a reasonable alternative within this environment, where coordination failure limits the usefulness of fully rational equilibrium concepts.

In their seminal work, McKelvey and Palfrey (1995) argue in favor of a more general interpretation of Quantal Response Equilibrium; namely, as a purely statistical model. While QRE is based upon a fixed-point argument, it may be thought of as merely a tool for organizing data, similar to a regression model. From this perspective, estimates of the precision parameter, λ , represent a relative measure of the degree of best response of subjects to the incentives of the game. It is reasonable to remain agnostic regarding the remaining, unexplained behavior. While the leading interpretation of this gap between random behavior (uniform mixing) and pure best-response is bounded rationality, conceptually, it is analogous to a statistical error term. Thus, a story of bounded rationality need

not be mutually exclusive with alternative interpretations of coordination failure, such as fairness or salience.

For example, one might imagine that the mixed-strategy equilibrium in which each player enters with probability $p = (c - 1)/(n - 1)$ is a focal point, since it is the only symmetric equilibrium in this anti-coordination game.³¹ However, the data presented here strongly reject this particular equilibrium. Alternatively, one might imagine that it is only fair to allow high types to enter more often than low types. After all, each subject is likely to have a turn to enter as a high type in future rounds. This is similar to the notion of payoff-dominance. QRE is perfectly consistent with either story, and the relative responsiveness of players to the incentives will determine the size of such effects.

However, this analysis raises a critical question. Why bother calculating this fixed-point, rather than simply estimating a conditional logit regression model of the style introduced by McFadden (1974)? Both should be capable of reasonably organizing the data. Both assume individuals will enter in proportion to expected payoffs. Both use the logit specification to transform expected payoffs into choice probabilities. Both are statistical models that may be fit to the data using maximum likelihood estimation techniques.

³¹Anti-coordination games are generally associated with asymmetry, which is why they are presumably harder to solve than standard coordination problems.

The answer is that the regression approach ignores strategic interaction. Small trembles by any individual affect the expected payoffs, and, hence, the choice probabilities of every other player. This feeds back into the original trembler's expected payoff, affecting the degree to which she trembles, and, again, affecting everyone else's payoffs, etc. QRE is interesting because of the discipline imposed by the fixed-point. Unlike a logit regression, QRE accounts for the feedback loop of noisy behavior on beliefs; which is potentially quite serious in an environment given the strong strategic inter-dependency of the coordination problem.

Table 2.9: Logit QRE versus Logit Regression

Group-size	Information	Simple	Fixed effect by Subject	Fixed effect by Round	Logit QRE	Logit Regression
Small	Complete	0.081*** (0.027)	0.105*** (0.025)	0.083*** (0.028)	0.104	0.121
Small	Incomplete	0.058** (0.028)	0.061** (0.025)	0.066** (0.028)	0.056	0.053
Large	Complete	0.035 (0.030)	0.023 (0.027)	0.038 (0.031)	0.047	0.092
Large	Incomplete	-0.015 (0.030)	-0.012 (0.027)	-0.024 (0.030)	0.042	0.076

Does this *reverberation effect* across players actually matter when trying to understand coordination failure in this game? The answer is clear: yes. As can be seen in Table 2.9, the estimated regression model predicts a higher entry rate for high types; however, it does not detect the attenuation of this effect as group-size increases. That is because this effect is driven by the combination of noisy best-response and strategic interaction, i.e. it is characteristic of the very strategic uncertainty that leads to coordination failure in the first place (van Huyck et al.,

1990). QRE gives economists a method for modeling of this strategic uncertainty and its effects. It is a useful model for this reason; regardless of the interpretation given to the departure from pure best-response.

2.5 Conclusion

There are natural extensions to this analysis. Experiment I presented toll effects, which effectively lowered the fixed income to entry, $a < r$. Meanwhile, experiment II varied the payoff magnitudes for individuals, v_i . In more recent work, Hartman (2010) has conducted an experiment including tolls ($a < r$) and various payoff magnitudes (v_i) simultaneously. The result is genuine preference heterogeneity. In fact, the toll induces a reduction in entry among low payoff types, allowing for sorting of high types onto the bridge. The fact that tolls solve the coordination problem by sorting the highest types onto the bridge speaks to its efficiency as a policy tool. It would be valuable to know how QRE would adjust to such preference heterogeneity.

From an econometric perspective, I could also allow for heterogeneity in the model by using heterogeneous QRE (Rogers et al., 2009). Allowing individuals to have unique noise parameters vastly expands the degrees of freedom of both QRE and the logit regression model, which may simple lead to over-fitting. That being said, differences may emerge between the two methods that lead to further

behavioral insights. For example, the estimated differences in the logit parameter across treatments found here may reflect the attenuation bias that arises from incorrectly specifying a homogeneous logit parameter (Golman, 2011). A recent development by Bajari and Hortacısu (2005) has lead to the estimation of HQRE without having to solve for a fixed point (Camerer et al., 2011).

Chapter 3

Time preferences and the management of coral reef fisheries

3.1 Introduction

Individuals' time preferences (here discounting and present bias) have been extensively researched as they pertain to demographic characteristics and financial decisions (Frederick et al., 2002). Recently, theories describing how individuals conceive of decisions and tradeoffs have begun to be applied more expansively, and research has considered the environmental implications of time preferences (Hardisty and Weber, 2009). Much of the environmental research to date has focused on how social discounting could influence policies for mitigating global warming (Carson and Roth Tran, 2009). Less research has focused on the ma-

rine realm, although notable exceptions include applications of time preference concepts to marine protected area design (Grafton et al., 2005; Sanchirico et al., 2006) and to ecosystem restoration (Sumaila, 2004). Research on the relationship between the discount factors³² of individuals and the management of marine resources is sparse.

Open access problems aside, we hypothesize that individuals with higher discount factors and less present bias with regard to financial decisions (i.e., those who value the future more highly) would also be more inclined towards resource conservation (i.e., marine reserves and less damaging types of fishing gear). Conversely, one might expect that individuals with lower discount factors and more present bias would be more inclined towards unsustainable levels of resource exploitation. Little empirical work has focused on this theory as pertains to fisheries management. Substantially more research has addressed the risk preferences of fishers (Bockstael and Opaluch, 1983; Eggert and Lokina, 2007; Eggert and Martinsson, 2004; Eggert and Tveteras, 2004; Mistiaen and Strand, 2000; Opaluch and Bockstael, 1984; Smith and Wilen, 2005) than the time preferences of fishers.

To our knowledge, only two published studies present fishers' discount factors. Both of those studies elicited individual discount factors (IDFs) using hypothetical choices between various fisheries management regimes and the theoretical fu-

³²Throughout the paper, we refer to discount factors, not discount rates. For clarity, if ρ is the discount rate and δ is the discount factor, the two are related by the equation $\delta = (1 + \rho)^{-1}$.

ture income streams associated with those regimes (Akapalu, 2008; Curtis, 2002). There do not appear to be any published studies presenting time preferences elicited from SCUBA divers. Thus, the research presented here represents the first attempt to elicit fishers' and divers' time preferences using incentivized experiments (i.e., price lists associated with actual monetary payments), and further, to explore the relationships between experimentally-measured discount factors and stated resource management preferences.

We elicited time preferences from fishers and professional SCUBA divers on Curaçao and Bonaire, islands in the southeastern Caribbean. Those professions were targeted because both are financially dependent on the health of ocean resources - fishers for the abundance of their catches, and professional divers for attracting tourist clientele. These neighboring islands are former Dutch colonies with similar histories of resource exploitation and similar marine ecosystems. The time preference experiment was paired with a socioeconomic interview that included questions on fishing and diving practices, perceptions of fish population trends and coral reef health, and level of support for management options such as gear restrictions and marine reserves. Here, we evaluate time preferences, as well as demographic characteristics, to understand fishers' and divers' preferred strategies for managing coral reefs.

3.2 Methods

3.2.1 Socioeconomic interviews

In fall of 2009 on Curaçao, and spring of 2010 on Bonaire, A.E.J. conducted in-person interviews with (full and part-time) fishers and professional SCUBA divers (i.e., dive instructors and divemasters). There are no records listing the fishers or divers on either island, so stratified random sampling of these groups was not possible. Instead, interviews were opportunistic, as exhaustive as possible, and as inclusive as possible of all demographic groups. Interviewees were identified via recommendations from local contacts, approaching individuals at fishing docks and in dive shops, and requesting the contact information for additional individuals at the end of each interview in what is termed a snowball sampling technique (Bernard, 1994). All divers were fluent or nearly fluent in English, and a Papiamento-Dutch-English translator was used for all fisher interviews.

A total of 388 interviews were conducted: 126 fishers on Curaçao, 51 fishers on Bonaire, 112 divers on Curaçao, and 99 divers on Bonaire. Based on the number interviews, the number of potential interviewees identified but with whom it was not possible to schedule interviews, and general knowledge of the fishing and diving communities, we estimate there are approximately 200 fishers on Curaçao, 80 fishers on Bonaire, 120 professional divers on Curaçao, and 130 professional

divers on Bonaire as of 2010. Based on these estimates, our sample represented 63% and 65% of the fishers on Curaçao and Bonaire respectively, and 86% and 83% of the divers on Curaçao and Bonaire respectively.

Of interviewees, eight fishers and five divers declined to participate in the time preference experiment because they refused to have their participation in the interview be at all associated with a monetary payment. Nine fishers and eleven divers had multiple switch points in one or more price lists. Because such responses imply either that this is an inappropriate approach for measuring IDFs for those interviewees, or that they did not properly understand the questions, those individuals are not included in this analysis. Five fishers did not provide full demographic information. Thus, the responses of 153 fishers and 197 divers are examined here.

3.2.2 Eliciting Time Preferences

Methods for eliciting time preferences have become well-honed, and the research presented here utilizes the best techniques currently available in attempt to capture the most accurate responses (Coller and Williams, 1999). Price lists, sets of questions offering choices between receiving payments sooner and later, accompanied by real monetary payments were used to elicit time preferences. At the end of each socioeconomic interview, participants were asked twenty-one ques-

tions - three price lists were used, each with seven questions (Appendix A). All price lists presented choices between sooner, smaller payments, and later, larger payments. Payments ranged from twenty to fifty florins (Fl.; 1 USD = Fl. 1.75). This maximum payment of Fl. 50 was chosen because it is roughly equivalent to a fisher or diver's daily income, thus one would not expect participants to be indifferent between payment choices. Additionally, it is a denomination of the local currency, so participants should have been familiar with its purchasing power, yet the amount is not so high as to make the experiment cost prohibitive.

The quantities of money offered were consistent across price lists. All sooner payments ranged from Fl. 50 down to Fl. 20, while all later payments were held constant at Fl. 50. The sole difference among the price lists was the dates at which payments were to be distributed. The first price list contained choices between payments the upcoming Friday and payments two weeks from Friday. The second price list contained choices between payments Friday and one month from Friday. The third price list contained choices between payments two weeks from Friday and a month from Friday. The experiment instructions and all questions were read aloud to interviewees.

To encourage careful consideration of responses, each interviewee was offered a cash payment in accordance with their answer to one of the twenty-one time preference questions. After responding to all questions, participants pick a num-

bered chip from a sack, and the quantity of their payment was determined based on how they answered the question corresponding with the number on the chip. For example, an interviewee who chose the chip marked with number seven, and who in response to question seven chose to receive Fl. 50 two weeks from Friday over Fl. 20 on Friday, would then actually be given Fl. 50 two weeks from Friday.

We employed front end delays (i.e., no payments were made at the time of interview) to equate the transaction costs of choosing the sooner and later payments, and to reduce the dependence of responses on level of trust for the researcher (Cardenas and Carpenter, 2008). All interviewees were required to retrieve their payments at a specified future date, time, and location. On Curaçao, we distributed payments on Friday afternoons at Dienst Landbouw, Veeteelt & Visserij (LVV), where the fisheries department is located. On Bonaire, we distributed payments on Friday afternoons at the office of Stichting Nationale Parken (STI-NAPA), the headquarters of the island's marine park. Each interviewee was given a card stating the date and time their payment could be picked up along with directions to the payment distribution location.

3.2.3 Calculating discount factors and present bias

For the point in each price list where the participant switched from preferring the sooner to preferring the later payment (i.e. the switch point), we took the mean

between the sooner payment amounts in the question before the switch and in the question where the switch was made. The mean is used, as is common practice, because price list questions do not enable determination of the exact switch point, rather the discrete range within which the switch occurs. We then divided that mean by Fl. 50, the highest payment option in each question, yielding discount factors from ≥ 1.0 down to 0.4 (Table 3.1). For example, a participant chooses sooner payments over later payments in response to all price list questions until asked to choose between Fl. 50 in two weeks from Friday and Fl. 20 on Friday, and at that point chooses to wait two weeks to receive Fl. 50. He would have a mean switch point of 25 (the mean of the Fl. 30 and Fl. 20 sooner payment amounts between which the switch was made), which when divided by 50 yields a discount factor of 0.5.

Table 3.1: Price list switch points with implied discount factors

Interviewee Preference	Midpoint between switch questions	Discount factor	# of fisher responses	# of diver responses
50 later over 50	50	≥ 1.0	31	58
50 sooner over 50 later	49	0.98	54	80
48 sooner over 50 later	46	0.92	10	16
44 sooner over 50 later	42	0.84	7	17
40 sooner over 50 later	37.5	0.75	20	13
35 sooner over 50 later	32.5	0.65	6	5
30 sooner over 50 later	25	0.5	7	6
20 sooner over 50 later	20	≤ 0.4	24	10

A participant who chose later payments in response to all questions, even when given the choice between Fl. 50 on Friday and Fl. 50 in 2 weeks, would

have a discount factor of 50/50, which we notate as ≥ 1.0 because it cannot be assigned an upper bound. A participant who chose sooner payments in response to all questions would have a discount factor of 20/50, which we notate as ≤ 0.4 because it cannot be assigned a lower bound. These bounds cannot be determined because the price lists did not include enough questions to determine switch points for these participants.

We consider time preferences using a three period $\beta - \delta$ model:

$$U_i(c_0, c_1, c_2) = u_i(c_0) + \beta\delta u_i(c_1) + \beta\delta^2 u_i(c_2) \quad (3.17)$$

In this model, $\beta = 1$ corresponds to the standard exponential discounting model, while $\beta \neq 1$ represents hyperbolic time preferences (i.e., changes in discount rates over time). Period zero refers to the upcoming Friday, while period one refers to two weeks from Friday, and period two refers to four weeks from Friday. Individual discount factors (IDF) were calculated based on participant responses to questions in the three price lists: IDF_1 , IDF_2 , and IDF_3 correspond with price lists one, two, and three, respectively. Within this context, our elicited discount factors relate to the model parameters as follows:

Discount Factors	Model Parameters
IDF_1	$\beta\delta$
IDF_2	$\beta\delta^2$
IDF_3	δ

Within the $\beta - \delta$ framework, IDF_3 yields the most direct measure of the discount factor, δ , so we exclusively use IDF_3 when analyzing the effect of discount factors on management preferences. To examine the role of hyperbolic discounting, individuals were categorized as present-biased, future-biased, or non-biased (pure exponential) based on the definitions from the underlying model. Specifically, we use the relationship between IDF_1 and IDF_3 to impute β , since these are the IDFs calculated from price lists with two-week differences in payment dates, but with IDF_3 having an additional two-week front-end delay.

Time Preferences	Value of β	Relation to IDFs
Present-biased	$\beta < 1$	$IDF_1 < IDF_3$
Exponential	$\beta = 1$	$IDF_1 = IDF_3$
Future-biased	$\beta > 1$	$IDF_1 > IDF_3$

Individuals for whom IDF_1 was smaller than IDF_3 are considered present-biased, since this implies hyperbolic discounting, i.e., $\beta < 1$. Individuals for whom IDF_1 was larger than IDF_3 are considered future-biased, i.e., $\beta > 1$. Individuals for whom IDF_1 is equal to IDF_3 are considered non-biased, corresponding to the standard model of exponential discounting.

3.2.4 Eliciting management preferences

As part of the larger interview, all participants were asked their opinions of a variety of marine resource management options (see Appendix B). Questions

were focused on fishing gear restrictions (i.e., requiring modifications to certain types of fishing gear or banning gear types all together) and area restrictions (i.e., temporary or permanent no fishing or no diving reserves). Marine reserves, where fishing is prohibited, are a scientifically proven way to increase fish populations, restore habitat, and improve catches in surrounding areas (for Interdisciplinary Studies of Coastal Oceans, 2007).

Using the nine gear questions and seven reserve questions, gear scores and reserves scores were calculated for each individual based on the percentage of gear restrictions and area restrictions they supported. Using responses to gear and reserve questions and five additional miscellaneous management questions (i.e., restricting numbers of fishers and divers, and prohibitions on anchoring, catch of certain species, and catch of juvenile fish), an overall conservation score was calculated for each individual. The maximum value for each score is 100 (i.e., all responses favoring restrictions on use and increased resource protection) and the minimum value is zero.

3.2.5 Revealed management preferences

Questions on limiting the number of fishers and divers, closing areas to fishing or diving, and restrictions on types of fishing gear are used to explore interviewees willingness to have their own usage constrained, that is, their revealed manage-

ment preferences. The most commonly used types of fishing gear on these islands are hook-and-line, trolling, fish traps, spearguns, gill nets, and beach seines. Hook-and-line and trolling generally cause less environmental harm (e.g., habitat damage, bycatch of juveniles and non-target species) than do fish traps, spearguns, gills nets, and beach seines. Thus, we term the former low-impact gears, and the latter high-impact gears. Low-impact gears have high risk of zero catch, but also a chance of valuable catch; they tend to catch low numbers of high-value species. High-impact gears have low risk of zero catch, are often less time intensive per kilogram caught (traps and nets need not be constantly attended while they are fishing), and have greater catch quantity, but with catch often (with the notable exception of spearguns) comprised of lower-value species.

3.2.6 Data analysis

We conduct a variety of regressions in order to estimate the impact of time preferences on attitudes towards management. Our model for each management score (denoted *Score*) takes the form:

$$\begin{aligned} \text{Score} = & \alpha + \phi_1 \text{IDF}_3 + \phi_2 \text{IDF}_3 \cdot \text{Fisher} + \theta_1 \text{Present} + \theta_2 \text{Present} \cdot \text{Fisher} \\ & + \theta_3 \text{Present} \cdot \text{Curacao} + \theta_4 \text{Present} \cdot \text{Fisher} \cdot \text{Curacao} + x' \lambda + \varepsilon \end{aligned} \tag{3.18}$$

where IDF_3 measures the exponential discount factor, δ ; *Present* is a dummy variable denoting present-biased time preferences, $\beta < 1$; *Fisher* is a dummy variable denoting fishers; *Curacao* is a dummy variable denoting residents of Curaçao; and x is a vector of demographic control variables. Within x are additional dummy variables, several continuous variables, and a few controls for interactions between variables. The dummy variables are for profession, location, marriage status, parenthood, possession of a bank account, access to credit, and employment status (i.e., full-time, part-time, hobby). The continuous variables are for age, years of experience, and a quadratic relation for the number of generations one's family has fished/dived. Lastly, to add more flexibility to the model, we also control for interactions between profession and island, profession and marriage status, and profession and possession of bank account.

We employ one large model to allow the effects of discount factors and present bias to vary by group, and to avoid the loss of statistical power associated with reductions in sample size. We denote the effect of IDF as ϕ_1 for divers and $\phi_1 + \phi_2$ for fishers. Likewise, we denote the effect of present bias as θ_1 for divers on Bonaire, $\theta_1 + \theta_2$ for fishers on Bonaire, $\theta_1 + \theta_3$ for divers on Curaçao, and $\theta_1 + \theta_2 + \theta_3 + \theta_4$ for fishers on Curaçao. For robustness, we estimate a few alternative restrictions to this flexible model, such as applying a uniform effect of discount factors ($\phi_2 = 0$) or present bias ($\theta_2 = \theta_3 = \theta_4 = 0$) on management preferences, or allowing the

effect of present bias to vary by profession ($\theta_3 = \theta_4 = 0$), island ($\theta_2 = \theta_4 = 0$), or both ($\theta_4 = 0$). Under the full model, the effect of IDF varies by profession, while present bias has four, independent effects for each profession-island pairing.

3.2.7 Risk-aversion

Any experiment involving time preferences must take care to address uncertainty and risk aversion. Our use of a front-end delay for all payments controls for changes in risk preferences within an individual due to the certainty of payments in the current period versus the uncertainty of payments in future periods (Andreoni and Sprenger, 2012). However, theory posits that a longer length to maturity of any payment is inherently riskier, and the yield curve (i.e., the fact that interest rates tend to rise with the length to maturity) is often cited as an example of this effect. The future is deemed riskier because there is some increasing probability that the payment will not be received. This reflects increasing riskiness of the asset, not increasing risk aversion of the individual, over the length to maturity. Recent research on U.S. treasuries (Startz and Tsang, 2012) suggests that the majority of the yield curve is explained by pure hyperbolic discounting, not risk. That is, what appears as an aversion to increasing risk over time-to-maturity is, in fact, increasing impatience with longer waiting periods. Hence, it may be suf-

ficient to control for present bias, which we do, without also controlling for risk aversion.

Theory gives us further reason to believe that we have successfully identified the effect of present bias in particular. While heterogeneous risk preferences may affect observed discount factors (δ) and, hence, observed hyperbolic discount factors (β), it does not affect whether hyperbolic discounting is observed on the extensive margin, that is, whether our dummy variable *Present* equals zero or one. Additionally, the time to payment is quite short (four weeks at most), which explains the large proportion of discount factors observed between 0.98 and 1, and mitigates concern over payoff uncertainty. Thus, we are confident that our results identify individual discount factors and hyperbolic discounting separately from risk preferences.

3.3 Results and Discussion

3.3.1 Interviewee demographics

The mean age of interviewed fishers was 48.0 years, significantly older than divers, for whom the mean age was 36.6 ($p < 0.001$, Table 3.2). Fishers were more likely than divers to be married ($p = 0.005$) and to have children ($p < 0.001$). Of married interviewees, divers were more likely than fishers to have an employed

spouse ($p < 0.001$). Proportions of fishers and of divers who were married, had children, and/or had an employed spouse did not differ significantly between islands. Fishers had more years of personal experience and more generations of family experience with fishing than divers had with diving (both $p < 0.001$).

Divers were significantly more financially secure than fishers, as measured by their greater likelihood of having a bank account, a credit card, and a friend who would loan them Fl. 50. (all $p < 0.001$). Of participants who could borrow Fl. 50 from a friend, fishers were significantly more likely to have to pay interest on such a loan ($p < 0.001$). Professional divers frequently work in the dive industry to temporarily support themselves while living abroad, tend to be more educated than fishers, and often have higher-paying jobs they can return to and/or family members (parents or spouses) to provide financial assistance while they pursue professional diving. Mean annual incomes of interviewees were approximately Fl. 22,100 (\$12,486 USD), and not significantly different between professions or islands. However, because income data were self-reported and 125 interviewees (35%) chose not to respond to this question, income data are not used further in the present analysis to avoid potential effects such as sample selection bias.

Table 3.2: Summary of demographic factors and management scores

	Fishers			Divers		
	Curaçao	Bonaire	Overall	Curaçao	Bonaire	Overall
Island						
# of interviewees	109	44	153	106	91	197
Mean age	47.90 (1.23)	48.50 (2.89)	48.00 (1.20)	34.20 (0.89)	39.30 (1.31)	36.60 (0.79)
Years of experience	34.71 (1.24)	38.14 (2.74)	35.69 (1.19)	12.16 (0.83)	16.77 (1.18)	14.29 (0.72)
# of generations	3.14 (0.14)	3.41 (0.18)	3.22 (0.12)	1.34 (0.06)	1.37 (0.06)	1.36 (0.04)
Married	0.55 (0.05)	0.43 (0.08)	0.52 (0.04)	0.36 (0.05)	0.39 (0.05)	0.37 (0.03)
Children	0.80 (0.04)	0.71 (0.07)	0.77 (0.03)	0.26 (0.04)	0.26 (0.05)	0.26 (0.03)
Bank account	0.72 (0.04)	0.66 (0.07)	0.70 (0.04)	1.00 (0.00)	0.98 (0.02)	0.99 (0.01)
Credit card	0.50 (0.05)	0.41 (0.07)	0.48 (0.04)	0.66 (0.05)	0.70 (0.05)	0.68 (0.03)
IDF_1	0.88 (0.02)	0.78 (0.03)	0.85 (0.02)	0.92 (0.01)	0.99 (0.02)	0.91 (0.01)
IDF_2	0.83 (0.02)	0.69 (0.04)	0.79 (0.02)	0.90 (0.01)	0.89 (0.02)	0.90 (0.01)
IDF_3	0.86 (0.02)	0.71 (0.04)	0.82 (0.02)	0.92 (0.01)	0.89 (0.02)	0.91 (0.01)
Present bias	0.16 (0.03)	0.07 (0.03)	0.13 (0.03)	0.12 (0.03)	0.10 (0.03)	0.11 (0.02)
Reserve Score	28.90 (3.36)	52.04 (4.12)	35.77 (2.80)	64.80 (2.17)	74.65 (2.17)	69.35 (1.57)
Gear Score	51.92 (1.87)	57.24 (3.91)	53.45 (1.75)	75.00 (1.98)	75.63 (2.52)	75.29 (1.57)
Conservation Score	44.28 (1.53)	52.87 (3.18)	46.75 (1.45)	71.36 (1.36)	74.44 (1.68)	72.78 (1.07)

To summarize the demographics, professional divers are generally younger, unmarried, without children, white, foreign (mostly Dutch), and more financially secure, with less personal and family history in their field than fishers. In contrast, fishers are generally older, married, parents, black, Antillean, less financially secure, and more experience. Diver demographics likely remain fairly consistent over time despite high turnover within the professional diving community. Turnover of fishers is low, although the proportion of time an individual allocates to fishing of-

ten varies seasonally and inter-annually based upon the catch and the availability of alternative employment opportunities.


3.3.2 Fisher and diver IDF's

The majority of fishers and divers appear to have relatively high discount factors. However, fishers' discount factors were significantly lower than those of divers. Mean IDF_3 for fishers was 0.82 (SE 0.019) compared to 0.91 (SE 0.011) for divers (Table 3.2). Divers were significantly more likely than fishers to choose the later payment even when the sooner and later payment amounts were both Fl. 50, which results in some IDF's of ≥ 1.0 ($p < 0.05$ for all price lists; Figure 3.1). Fishers were significantly more likely than divers to choose the sooner payment for all questions in a price list ($p < 0.002$ for all price lists). For fishers, IDF_3 was positively correlated with having a bank account ($p = 0.013$). Counter to our hypothesis, IDF's were not correlated with usage of high-impact gear.

Although evidence in the literature for the correlation between poverty and time preferences is mixed (Cardenas and Carpenter, 2008), some research has identified a positive relationship between discount factor and income (Carson and Roth Tran, 2009). Thus, the results here could reflect the fact that divers generally have higher expected lifetime incomes than fishers, and a lower level of financial constraint.

Individuals largely offered similar responses across price lists. Fifteen percent of fishers and 27% of divers had the same switch point in all three price lists. Consistency in switch points could be caused by individuals not perceiving these lists as presenting substantially different choices. Larger differences in days until payment (time periods on the order of six months or a year) might have produced greater differentiation in responses between price lists, but could also have exacerbated any effect of risk preferences.

Figure 3.1: Distributions of individual discount factors



3.3.3 Distribution of IDFs

The majority of both fishers and divers have discount factors ≥ 0.98 , however, the distribution of diver IDFs exhibits an essentially monotonic decline from 1.0 down to ≤ 0.4 for all price lists, while that of fishers has a somewhat tri-modal distribution with additional peaks at 0.75, and at ≤ 0.4 (Figure 3.1). The cluster of fishers at $\text{IDF} \leq 0.4$ represents 9.7% of participants in the sample. This cluster may be an artifact of the truncation of each price at the choice between Fl. 20 sooner and Fl. 50 later. If the range of sooner payments had included amounts less than Fl. 20, some participants may have chosen additional lower payments, resulting in discount factors lower than 0.4, thus dispersing that peak.

The cluster of fishers at an IDF of 0.75 represents approximately 10% of participants. This concentration of responses, which corresponds with a preference of Fl. 40 sooner over Fl. 50 later, but a preference of Fl. 50 later over Fl. 35 sooner and, may be meaningful and not an artifact of the experimental design. Currency quantities are often mentally compared against the value of oft purchased items, offering a potential explanation for this peak. Polar, the most popular beer on both islands, was priced from Fl. 36 to Fl. 38 per case at the time of this study, and is one likely candidate for such a calibration. Although anecdotal and inad-

equate to prove correlation, five fishers and one diver did explicitly mention this as the rationale for their switch point.

3.3.4 Comparison with previous IDF studies

Despite the array of literature on measuring time preferences, and due in part to widely varying elicitation methods, honing in on a “normal” range of discount factors has been elusive. As compiled by Frederick et al. (2002), seven studies published between 1978 and 2002 elicited discount factors using choice sets associated with real monetary payments. Those studies produced a range of annual discount factors from 0.0 to 1.01. The range of IDFs presented here, from ≤ 0.4 to ≥ 1.0 , is truncated due to the experimental design, hindering direct comparison with previous research.

The two previous studies that report experimentally measured discount factors of fishers found mean values substantially lower than those presented here. Fishers in the Irish Sea had a mean discount factor of ~ 0.7 (Curtis, 2002), and fishing boat skippers in Ghana had a mean discount factor of 0.43 (Akapalu, 2008). These studies measured IDF with questions about future profit scenarios that would result from different fishery management approaches. Because those two studies used different elicitation methods than the price list approach used here, unfortunately it cannot be determined whether the higher discount factors we

measured are indicative of important differences in cultural, economic, or fisheries management contexts.

Another confounding factor when comparing discount rates is time period. Broadly, there is the question of determining the appropriate time horizon for considering the relationship of time preferences with ocean management decisions, given that costs are often incurred in the near-term, while benefits accrue at an unknown rate in the future. That said, we found sufficient variation in discount factors within our two-to-four week time period to explain management preferences (Table 3.3). The weak incentives implied by a short time horizon should only make significant results harder to find. Further, longer time periods to payment increase the potential for participants to factor uncertainty of payment into their decision-making, which could confound time preferences. Hence, we are satisfied that the time period we used is sufficient to meaningfully explain management preferences.

3.3.5 Present bias

Distributions of hyperbolic discounting were not significantly different between professions or islands (Table 3.2). Contrary to our hypothesis, fishers who use high-impact gear did not exhibit a higher incidence of present bias. Two-thirds of both fishers and divers were non-biased. An additional 22% of participants were

future-biased, and the remaining 12% were present-biased. As with IDF_3 (Section 3.2), the relatively low occurrence of present bias may represent consistency across price lists that individuals perceived as functionally equivalent. As for the large proportion of individuals exhibiting future bias, perhaps individuals were using the experimental payment as a commitment device for saving. For example, one interviewee explained that although he did not currently have a girlfriend, he might in two weeks, and if he waited for the payment then he would have money to take her out to dinner.

Alternatively, motivation to wait for payment could be related to the highly variable nature of fisher and diver incomes, with fishers depending on unpredictable catches and divers depending partly on tips. Future bias could be related to whether the interviewee had recently received a paycheck, whether divers had recently received a good tip, and whether fishers recently had a good catch. Interviewees may also have been hesitant to admit a preference for receiving money sooner, due to pride. This hesitation could have been exacerbated by the fact that the participants were mostly male and the interviewer was a young female. Indeed, decisions can be influenced by socio-cultural context, the desire to appear competent, and the interview process itself, and therefore may not solely reflect monetary preferences (Bowles, 1998; Levitt and List, 2007). Regardless, future bias has no measurable impact on management scores for any of the functional

forms we estimate. Thus, it is included within the reference category, along with exponential discounting, in order to increase the statistical power of our measurement of the effect of present bias.

3.3.6 Management scores

Mean gear scores, reserve scores, and conservation scores were significantly higher for divers than for fishers (all $p < 0.001$, Table 3.2). This is not surprising, since divers benefit from restrictions on fishing, while fishers bear the bulk of costs for such restrictions. Without conditioning on any control variables, IDF_3 (Table 3.3) is positively correlated with reserve scores (R1, $p = 0.068$), gear scores (G1, $p = 0.023$), and conservation scores (C1, $p = 0.003$). This is consistent with our hypothesis that more patient individuals would show greater support for conservation efforts. Yet, when considering fishers and divers collectively, after controlling for other relevant factors by estimating our model with a uniform effect of discount factors and present bias, there are no significant relationships between these scores and individual discount factors (R2, G2, C2).

Table 3.3: Marginal effect of discount factors on management scores

Variable	Restriction	All Divers	All Fishers
Reserve Score	R1: $\phi_2 = \theta_1 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	17.37* (9.50)	17.37* (9.50)
	R2: $\phi_2 = \theta_2 = \theta_3 = \theta_4 = 0$	7.09 (9.16)	7.09 (9.16)
	R3: $\theta_2 = \theta_3 = \theta_4 = 0$	25.30*** (9.39)	-6.31 (13.99)
	R4: Full Model	24.81*** (9.52)	-6.43 (14.17)
Gear Score	G1: $\phi_2 = \theta_1 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	14.87** (6.51)	14.87** (6.51)
	G2: $\phi_2 = \theta_2 = \theta_3 = \theta_4 = 0$	3.05 (6.55)	3.05 (6.55)
	G3: $\theta_2 = \theta_3 = \theta_4 = 0$	15.09 (11.84)	-5.82 (7.53)
	G4: Full Model	15.09 (11.93)	-6.45 (7.57)
Conservation Score	C1: $\phi_2 = \theta_1 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	15.61*** (5.29)	15.61*** (5.29)
	C2: $\phi_2 = \theta_2 = \theta_3 = \theta_4 = 0$	4.06 (4.84)	4.06 (4.84)
	C3: $\theta_2 = \theta_3 = \theta_4 = 0$	18.30** (7.26)	-6.43 (6.29)
	C4: Full Model	18.05** (7.33)	-6.70 (6.33)

However, if we allow discount factors to have different effects for fishers and divers, but present bias is still restricted to one overall effect, for fishers there is still no effect, but we see that for divers (R3, G3, C3) discount factors are significant in explaining reserve scores ($p = 0.007$) and conservation scores ($p = 0.012$), though not gear scores. A 0.1 point increase in the discount factor leads to a 2.5 point increase in reserve scores for divers, or an almost 15 point increase over the range of observed discount factors. Similarly, for conservation scores a 0.1 increase in the discount factor leads to a 1.8 point increase in conservation scores for divers, or 10.8 points over the observed range. These results are robust

to the restrictions we place on the effect of present bias (R4, G4, C4). Since conservation scores are composed in part from reserve scores and gear scores, the smaller effect of discount factors on conservation scores is expected. The lack of an effect of discount factors on gear scores is also expected, since marine reserve restrictions directly affect inter-temporal decision-making while gear restrictions have only second-order inter-temporal effects (e.g., restricting high-impact gear leads to improved future catches).

Table 3.4: Marginal effect of present bias on marine reserve scores

Restriction	Bonaire Divers	Bonaire Fishers	Curaçao Divers	Curaçao Fishers
P1: $\phi_1 = \phi_2 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	-13.92*** (5.26)	-13.92*** (5.26)	-13.92*** (5.26)	-13.92*** (5.26)
P2: $\theta_2 = \theta_3 = \theta_4 = 0$	-7.43** (3.60)	-7.43** (3.60)	-7.43** (3.60)	-7.43** (3.60)
P3: $\theta_3 = \theta_4 = 0$	-5.55 (4.24)	-9.71 (6.06)	-5.55 (4.24)	-9.71 (6.06)
P4: $\theta_2 = \theta_4 = 0$	0.26 (5.52)	0.26 (5.52)	-10.73** (4.42)	-10.73** (4.42)
P5: $\theta_4 = 0$	0.63 (5.97)	-0.88 (7.46)	-9.91** (5.02)	-11.43* (6.18)
P6: Full Model	-2.12 (6.69)	7.17 (8.02)	-7.93 (5.44)	-13.11* (6.77)

Present bias (Table 3.4) reduces marine reserve scores by almost 14 points (P1, $p = 0.068$), 7.4 points (P2, $p = 0.040$) after controlling for demographic variables, but does not significantly reduce gear scores or conservation scores. That is, in support of our hypothesis, interviewees who were more present biased were less supportive of marine reserves. Gear scores are not responsive to present bias, thus washing out an effect for overall conservation scores, even after controlling

for demographic variables. The effect of present bias on reserve scores is stronger for fishers than divers (P3, 9.7 versus 5.5), though not significant for either, and stronger on Curaçao than Bonaire (P4, 10.7 versus 0.26, $p = 0.016$). Model P5 estimates the effect of present bias across professions and islands simultaneously, though the model assumes the difference between fishers and divers is the same for each island. The estimates suggest a combined effect that is significant for both fishers ($p = 0.049$) and divers ($p = 0.065$) on Curaçao; though present bias appears irrelevant on Bonaire. The final model (P6), which estimates a unique effect of present bias for each profession-island combination, suggests that present bias only applies to fishers on Curaçao ($p = 0.054$). Given the relatively low statistical power of these estimates (four parameters with 40 observations of present bias), it is not clear whether the effect applies only to Curaçao fishers, or whether we lack the sample size required to parse the effect of present bias across profession and location at the same time. Perhaps these inter-island differences could be explained by the greater industrialization and faster pace of life on Curaçao relative to Bonaire (Wang et al., 2009).

3.3.7 Principal component analysis of management scores

The design of the management scores was carefully tailored to best measure the conservation preferences of participants. That said, these key variables are

constructed from survey responses; they are not naturally occurring and passively observable. This is true not only of conservation scores, which are composites of other created variables, but also, more subtly, of reserve scores and gear scores. Therefore, as a robustness check, we conduct a principal component analysis (PCA) on the management scores and repeat our regression analysis using the principal components as dependent variables.

We find that the effects of discount factors and present bias operate across distinct principal components. This result indicates that IDF_3 and *Present* are, indeed, measuring distinct aspects of time preferences. The first principal component is found to explain 76.9% of the variation across management scores, while the second and third components explain 20.8% and 2.3% of the remaining variation, respectively. While the units of the principal components, themselves, are not interpretable, the relative size of the estimated effects makes for useful comparisons.

The PCA results generally confirm our other findings. Only the first principal component is significantly affected by discount factors (Table 3.5). When allowing the effect of discount factors to vary across fishers and divers, we find that discount factors only impact the conservation attitudes of divers (PCA-R3, PCA-R4). In section 3.6, we showed discount factors impact both reserve scores and conservation scores for divers. Therefore, it appears the measured effect for

conservation scores merely represented the residual component of conservation scores attributed to reserve scores.

Table 3.5: Marginal effect of discount factors on the principal components

Variable	Restriction	All Divers	All Fishers
Reserve Score	PCA-R1: $\phi_2 = \theta_1 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	1.11*** (0.40)	1.11*** (0.40)
	PCA-R2: $\phi_2 = \theta_2 = \theta_3 = \theta_4 = 0$	0.31 (0.35)	0.31 (0.35)
	PCA-R3: $\theta_2 = \theta_3 = \theta_4 = 0$	1.31** (0.53)	-0.43 (0.46)
	PCA-R4: Full Model	1.30** (0.53)	-0.46 (0.46)
Gear Score	PCA-G1: $\phi_2 = \theta_1 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	-0.02 (0.24)	-0.02 (0.24)
	PCA-G2: $\phi_2 = \theta_2 = \theta_3 = \theta_4 = 0$	0.08 (0.27)	0.08 (0.27)
	PCA-G3: $\theta_2 = \theta_3 = \theta_4 = 0$	0.16 (0.39)	0.02 (0.38)
	PCA-G4: Full Model	0.15 (0.39)	0.03 (0.39)
Conservation Score	PCA-C1: $\phi_2 = \theta_1 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	-0.05 (0.07)	-0.05 (0.07)
	PCA-C2: $\phi_2 = \theta_2 = \theta_3 = \theta_4 = 0$	-0.01 (0.07)	-0.01 (0.07)
	PCA-C3: $\theta_2 = \theta_3 = \theta_4 = 0$	-0.06 (0.10)	0.04 (0.11)
	PCA-C4: Full Model	-0.06 (0.10)	0.03 (0.11)

The third principal component is the only dimension significantly affected by present bias (Table 3.6). Since the third component is the weakest, this explains the difficulty in finding consistent results for the effect on reserve scores. By isolating this component, we obtain stronger results for present bias. Varying the effect of present bias by profession (PCA-P3), this component is reduced by 0.11 for divers ($p = 0.028$). Allowing the effect to differ across islands (PCA-P4), we estimate the reduction to be 0.12 on Bonaire ($p = 0.016$) and 0.10 on

Curaçao ($p = 0.045$). Allowing each profession-island to have a unique present bias effect, we see that Curaçao fishers' scores are reduced by 0.13 ($p < 0.001$). Surprisingly, with PCA we find that Bonaire divers' scores are reduced by 0.18 ($p < 0.001$), a far larger reduction than before. Thus, the PCA shows that our previous estimates lacked the statistical power necessary to separate out the effects of present bias on reserve scores across profession and island. We have confirmed a strong positive effect of discount factors, and a significant negative effect of present bias, on conservation attitudes. These effects are measured across different principal components, so we are confident that the elicitation experiment separately identifies these two dimensions of time preference.

Table 3.6: Marginal effect of present bias on the third principal component

Restriction	Bonaire Divers	Bonaire Fishers	Curaçao Divers	Curaçao Fishers
PCA-P1: $\phi_1 = \phi_2 = \theta_2 = \theta_3 = \theta_4 = \lambda = 0$	-0.12*** (0.01)	-0.12*** (0.01)	-0.12*** (0.01)	-0.12*** (0.01)
PCA-P2: $\theta_2 = \theta_3 = \theta_4 = 0$	-0.10** (0.04)	-0.10** (0.04)	-0.10** (0.04)	-0.10** (0.04)
PCA-P3: $\theta_3 = \theta_4 = 0$	-0.11** (0.05)	-0.10 (0.59)	-0.11** (0.05)	-0.10 (0.59)
PCA-P4: $\theta_2 = \theta_4 = 0$	-0.12** (0.05)	-0.12** (0.05)	-0.10** (0.05)	-0.10** (0.05)
PCA-P5: $\theta_4 = 0$	-0.13** (0.05)	-0.12 (0.09)	-0.10 (0.07)	-0.09 (0.06)
PCA-P6: Full Model	-0.18*** (0.05)	0.05 (0.10)	-0.06 (0.08)	-0.13** (0.06)

3.3.8 Restricting their own resource use

We analyzed fishers' and divers' comparative willingness to put restrictions on their own resource use. One percent of fishers were willing to limit the number of fishers, whereas 34% of divers were willing to limit the number of divers ($p < 0.001$).³³ Fishers who used high-impact gears (traps, spearguns, gill nets, seines) had significantly lower marine reserve scores ($p = 0.044$), gear scores ($p = 0.014$), and conservation scores ($p = 0.037$) than those who use low-impact gears (hook and line, troll), reflecting reluctance not just on having restrictions placed on their use of these gears, but on conservation measures in general.

Fishers were more likely than divers to support no diving areas and vice versa for no fishing areas (both $p < 0.001$). Sixty-two percent of divers supported no diving areas versus 37% of fishers who supported no fishing areas. On Bonaire, the same relationship holds as regards the existing no fishing and no diving areas, with fishers more likely than divers to support additional no diving areas ($p = 0.035$) and vice versa for no additional fishing areas ($p = 0.008$). Divers on Bonaire were significantly more likely to support additional no diving areas (51% support) than fishers were to support additional no fishing areas (32% support; $p < 0.05$). Taken together, this is evidence of divers' greater patience with resource

³³A restriction on the number of divers is likely to restrict the number of tourist divers, and thereby the number of potential customers, but not necessarily the number of professional divers. A restriction on the number of fishers could directly prevent interviewees from fishing.

use relative to fishers. To be fair, we are comparing a non-extractive, recent profession dominated by foreigners (diving) with a resource extractive profession with a long cultural history conducted by locals who feel displaced (fishing). There are likely complex socio-cultural factors at play here. Regardless, divers focus on observing marine life and therefore benefit from restrictions on fishing, so these differences in observed management preferences are to be expected to some degree.

3.3.9 Key findings

Fishers and divers had relatively high discount factors; however, fishers' IDFs were significantly lower. The distribution of fishers' IDFs is somewhat trimodal (versus a monotonic decline for divers), perhaps related to the salient purchasing power of particular dollar amounts. Interestingly, divers' discount factors were positively related to their view of marine reserves, though not with their views of gear restrictions. Discount factors appear unrelated to management scores for fishers.

Fishers who use high-impact gear did not exhibit the higher incidence of present bias we had anticipated. However, as expected, interviewees who were more present biased were less supportive of marine reserves, controlling for a host of demographic and economic factors. An explanation for the high proportion of

present-biased interviewees cannot be determined from our data, but could be due to income variability or socio-cultural factors.

Divers are overall more supportive of management measures than fishers, and, divers' responses support our hypothesis that individuals who are more patient with money are more likely to support marine reserves. Also as expected, fishers who use high-impact gears are less supportive of marine reserves. Interestingly, divers on Bonaire, where there is a long history of marine protection, were the most supportive of conservation measures.

3.3.10 Policy implications

Fishers' relatively lower levels of financial patience and support of management measures that would contribute to long-term resource conservation has implications for the effectiveness of various policy approaches. Perhaps management of fishing and diving should be approached differently, with consideration of these inter-group differences in time preferences.

Our research suggests two options for applying our behavioral economic results for increasing fishery sustainability. First, policy could attempt to shift fishers' incentives towards conservation - perhaps with transfer payments to offset the near-term expenses of switching to low-impact gears, reducing fishing effort, or adding marine reserves. Since dive industry employees are more patient financially,

have greater access to credit, exhibit higher valuations of conservation, and would benefit from constraints on fishing, then perhaps they, or the dive industry more broadly, could offset the costs of altering fishing behaviors. Such an offset may be particularly needed because fishers who use high-impact gears are less supportive of conservation measures ($p = 0.037$).

The notion of offsets is timely in the context of the dive fee that exists on Bonaire and is being considered on Curaçao. Every individual who dives or snorkels on Bonaire is required to buy an annual marine park tag for 25 USD, the funds from which are used for marine park management. The idea for Curaçao is similar; an annual fee would be used to fund enforcement of marine regulations. Utilizing a portion of these funds to pay fishers to reduce or cease use of high-impact gears could be a cost-effective conservation approach. For example, this might be a mechanism for a buyout of the fish traps and nets that inflict damage on shallow reef ecosystems, primarily catch herbivorous fish (i.e. parrotfish and surgeonfish) critical for controlling algal growth, and are deeply disliked by divers. Buyouts or other offsets, however, would require substantial community buy-in and enforcement capacity.

Regarding community buy-in, there is an effect we call conservation inertia. In addition to Bonaire having a dive fee, it also has no fishing and no diving areas, whereas Curaçao has none of these. These differences seem to be associated with a

dramatic inter-island difference in level of support for marine reserves. No fishing areas receive 52% support on Curaçao versus 84% on Bonaire, and no diving areas receive 32% support on Curaçao versus 94% on Bonaire (both $p < 0.001$). Perhaps out of familiarity, people on Bonaire are more supportive of reserves, whereas people on Curaçao may resist reserve establishment because the consequences are unknown. Thus, gaining community support (especially from fishers) to put initial conservation measures in place may be a challenge, but conservationist views may grow after measures are put in place and benefits are seen.

Second, establishing property rights may not alone be sufficient for conservation. Ownership is presumed to lead to improved resource stewardship, and in recent years fisheries economists have examined the potential for property rights-based management approaches such as cooperatives, individual transferable quotas (ITQs), and territorial use rights in fisheries (TURFs). However, present bias can lead to over-consumption of resources in the present despite the incentive for sustainable use that is associated with ownership. For example, work by Hepburn et al. (2010) demonstrates how mechanisms that align incentives with resource-extraction externalities could still (though unlikely) leave the resource stock susceptible to inadvertent collapse in the presence of hyperbolic discounting. Therefore, even a property rights-based management scheme that addresses the common-pool resource problem may need to be paired with additional mea-

asures in order to prevent overfishing, if the fishers themselves are present-biased. Interestingly, fishers who more greatly discount the future have been shown to be more likely to violate fisheries regulations, especially if those regulations aren't perceived to be legitimate and the risk and severity of punishment are low (Akapalu, 2008).

This is not to say that management regimes involving property rights cannot be effective. In fact there is much evidence that the sustainability of fisheries can be improved by both individual transferable quotas (Costello et al., 2008) and territorial use rights fisheries (White and Costello, 2011), given certain ecological conditions and institutional structures. Thus, property rights should be implemented within an overall framework for sustainably managing resource use. When resource users are highly present-biased and financially constrained, the odds are stacked against sustainable use. Thus, to lower the chances of overexploitation, property rights could be paired with some mix of restrictions on effort and gears, near-term incentives for sustainable use, strong enforcement, and robust community buy-in. We suggest that future research into the most effective combinations of management tools would be an important contribution to the field.

Ineffective management of coral reef fisheries can result in overfishing, overcapitalization, and low profits (Munro and Scott, 1985), can facilitate a shift to an algal-dominated ecological state (Hughes, 1994; Knowlton, 2004), and can jeopard-

dize food security (Pauly, 2006; Sadovy, 2006). It is therefore critical for coral reef management efforts and for fisheries research to consider how to address the myriad factors that influence individuals' resource use patterns, and apply economic research when endeavoring to better align financial incentives with sustainability. There is an important role for behavioral economics to play in developing strategies to sustainably manage ocean resources.

Conclusion

This dissertation has made use of behavioral models, experimental data, and econometric techniques to gain a deeper understanding individual decision-making in three public goods problems. In so doing, I am able to accurately predict many useful comparative statics, as well as increase predictive power.

The first two chapters analyze the noisy behavior of subjects in two coordination/anti-coordination problems. Standard tools are indeterminate in the threshold game due to the multiplicity of equilibria, while treatment effects in the market entry game are designed to be neutral. In both cases, treatment effects emerge that cry out for an explanation. Given the noisiness of individual choices in these games, and given the failure of group behavior to ever converge, quantal response equilibrium represents a reasonable alternative. I find that structural estimates of the precision parameter yield useful predictions across many treatment effects.

The third chapter uses a simple $\beta\delta$ -model to motivate the construction of important variables in the data. Time preferences are elicited using a paid exper-

Conclusion

iment, and demographic data is collected in the survey portion of the interview. Controlling for key demographic factors, my co-author and I find that decreased patience (lower δ) and present bias ($\beta < 1$) reduce willingness to use marine reserves to conserve fisheries, but has no effect preferences for gear restrictions. This is sensible because marine reserves represent a fundamentally inter-temporal constraint on extraction, while gear restrictions are intended to mitigate a contemporaneous externality levied by fishers.

Various extensions lie ahead of this work. First, I would like to jointly estimate a logit QRE and a warm-glow parameter. Second, I look forward to working with John Hartman in integrating all three of our experiments, so that we may examine the effects of preference heterogeneity, not just payoff heterogeneity. Such extensions certainly push the boundary of current computational feasibility. This makes them challenging but, as yet, untapped lines of inquiry.

Bibliography

Bibliography

- Akapalu, W.**, “Fishing regulations, individual discount rate, and the fisherman behaviour in a developing country fishery,” *Environmental and Development Economics*, 2008, *13*, 591–606.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt**, “A theoretical analysis of altruism and decision error in public goods games,” *Journal of Public Economics*, 1998, *70*, 297–323.
- , – , and – , “Noisy Directional Learning and the Logit Equilibrium,” *Scandinavian Journal of Economics*, 2004, *106* (3), 581–602.
- Andreoni, J. and C. Sprenger**, “Risk preferences are not time preferences,” *The American Economic Review*, 2012, *102*, 3357–3376.
- Andreoni, James**, “Cooperation in public-goods experiments: Kindness or confusion?,” *The American Economic Review*, 1995, *85* (4), 891–904.
- and **John Miller**, “Giving according to GARP: An experimental test of the consistency of preferences for altruism,” *Econometrica*, 2002, *70* (2), 737–753.
- Arnott, Richard and Kenneth Small**, “The economics of traffic congestion,” *American Scientist*, 1994, *82*, 446–455.
- Bagnoli, Mark and Barton L. Lipman**, “Provision of Public Goods: Fully Implementing the Core through Private Contributions,” *The Review of Economic Studies*, 1989, *56* (4), 583–601.
- and **Michael McKee**, “Voluntary Contribution Games: Efficient Private Provision of Public Goods,” *Economic Inquiry*, 1991, *29* (2), 351–366.
- , **Shaul Ben-David, and Michael McKee**, “Voluntary provision of public goods: the multiple unit case,” *Journal of Public Economics*, 1992, *47* (2), 85–106.

BIBLIOGRAPHY

- Bajari, Patrick and Ali Hortacısu**, “Are structural estimates of auction models reasonable? Evidence from experimental data,” *Journal of Political Economy*, 2005, 113 (4), 703–741.
- Bernard, H.R.**, *Research Methods in anthropology: Qualitative and quantitative approaches*, Sage Publications, 1994. Thousand Oaks, CA, p.585.
- Bockstael, N.E. and J.J. Opaluch**, “Discrete modelling of supply response under uncertainty: The case of the fishery,” *Journal of Environmental Economics and Management*, 1983, 10, 125–137.
- Bowles, S.**, “Endogenous preferences: The cultural consequences of markets and other economic institutions,” *Journal of Economic Literature*, 1998, 36, 75–111.
- Camerer, Colin and Dan Lovallo**, “Overconfidence and Excess Entry: An Experimental Approach,” *The American Economic Review*, 1999, 89 (1), 306–318.
- , **Salvatore Nunnari, and Thomas R. Palfrey**, “Quantal response and nonequilibrium beliefs explain overbidding in maximum-value actions,” *Working Paper #1349*, 2011. California Institute of Technology.
- Cardenas, J.C. and J. Carpenter**, “Behavioural development economics: Lessons from field labs in the developing world,” *Journal of Development Economics*, 2008, 44, 311–338.
- Carson, R.T. and B. Roth Tran**, “Discounting behavior and environmental economics,” *Journal of Neuroscience, Psychology, and Economics*, 2009, 2, 112–130.
- Chen, Hsiao-Chi, James W. Friedman, and Jacques-Fracois Thisse**, “Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach,” *Games and Economic Behavior*, 1997, 18, 32–54.
- Coller, M. and M.B. Williams**, “Eliciting individual discount rates,” *Experimental Economics*, 1999, 2, 107–127.
- Costello, C., S.D. Gaines, and J. Lynham**, “Can catch shares prevent fisheries collapse?,” *Science*, 2008, 321, 1678–1681.
- Croson, Rachel T. A. and Melanie B. Marks**, “Step Returns in Threshold Public Goods: A Meta- and Experimental Analysis,” *Experimental Economics*, 2000, 2 (3), 239–259.

BIBLIOGRAPHY

- Curtis, J.**, “Estimates of fishermen’s personal discount rate,” *Applied Economics Letters*, 2002, 9, 775–778.
- Davis, Douglas D. and Charles A. Holt**, *Experimental Economics*, Princeton University Press, 1993.
- Duffy, John and Ed Hopkins**, “Learning, Information, and Sorting in market Entry Games: Theory and Evidence,” *Games and Economic Behavior*, 2005, 51, 31–62.
- Eggert, H. and P. Martinsson**, “Are commercial fishers risk-lovers?,” *Land Economics*, 2004, 80, 550–560.
- **and R. Tveteras**, “Stochastic production and heterogeneous risk preferences: Commercial fishers’ gear choices,” *American Journal of Agricultural Economics*, 2004, 86, 199–212.
- **and R.B. Lokina**, “Small-scale fishermen and risk preferences,” *Marine Resource Economics*, 2007, 22, 49–67.
- Erev, Ido and Amnon Rapoport**, “Coordination, “Magic,” and Reinforcement Learning in a Market Entry Game,” *Games and Economic Behavior*, 1998, 23, 146–175.
- Eyster, Erik and Matthew Rabin**, “Cursed equilibrium,” *Econometrica*, 2005, 73 (5), 1623–1672.
- Fischbacher, Urs**, “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 2007, 10 (2), 171–178.
- for Interdisciplinary Studies of Coastal Oceans, Paternership**, “The science of marine reserves (2nd edition),” 2007. www.piscoweb.org.
- Frederick, S., G. Loewenstein, and T. O’Donoghue**, “Time discounting and time preferences: A critical review,” *Journal of Economic Literature*, 2002, 40, 351–401.
- Gigerenzer, G. and D. Goldstein**, “Reasoning the Fast and Frugal Way: Models of Bounded Rationality,” *Psychological Review*, 1996, 103 (4), 650–669.
- Goeree, Jacob K., Charles A. Holt, and Angela K. Moore**, “An Experimental Examination of the Volunteer’s Dilemma,” *Working Paper*, 2005.

BIBLIOGRAPHY

- , – , and **Thomas R. Palfrey**, “Regular Quantal Response Equilibrium,” *Experimental Economics*, 2005, 8, 347–367.
- Golman, Russell**, “Quantal Response Equilibria with Heterogeneous Agents,” *Journal of Economic Theory*, 2011, 146, 2013–2028.
- Grafton, R.Q., T. Kompas, and D. Lindenmayer**, “Marine reserves with ecological uncertainty,” *Bulletin of Mathematical Biology*, 2005, 67, 957–971.
- Haile, Philip, Ali Hortaçsu, and Grigory Kosenok**, “On the empirical content of quantal response equilibrium,” *American Economic Review*, 2008, 91, 180–200.
- Hardisty, D.J. and E.U. Weber**, “Discounting future green: Money versus the environment,” *Journal of Experimental Psychology: General*, 2009, 138, 329–340.
- Hartman, John**, “A route choice experiment with an efficient toll,” *Working Paper*, 2007. Department of Economics, University of California Santa Barbara.
- , “The relevance of heterogeneity in a congested route network with tolls: An analysis of two experiments using actual waiting times and monetized time costs,” *Working Paper*, 2010. Department of Economics, University of California Santa Barbara.
- Hepburn, C., S. Duncan, and Papachristodoulou**, “Behavioral economics, hyperbolic discounting and environmental policy,” *Environmental and Resource Economics*, 2010, 46, 189–206.
- Houser, Daniel and Robert Kurzban**, “Revisiting kindness and confusion in public goods experiments,” *The American Economic Review*, 2002, 92 (4), 1062–1069.
- Hughes, T.P.**, “Catastrophes, phase-shifts, and large-scale degradation of the Caribbean coral reef,” *Science*, 1994, 265, 1547–1551.
- Isaac, Mark and James M. Walker**, “Group size effects in public goods provision: the voluntary contributions mechanism,” *The Quarterly Journal of Economics*, 1988, 103 (1), 179–199.
- Kagel, John H. and Alvin E. Roth**, *The Handbook of Experimental Economics*, Princeton University Press, 1997.

BIBLIOGRAPHY

- Knowlton, N.**, “Multiple “stable” states and the conversation of marine ecosystems,” *Progress in Oceanography*, 2004, *60*, 387–396.
- Levitt, S.D. and J.A. List**, “What do laboratory experiments measuring social preferences reveal about the real world?,” *Journal of Economic Perspectives*, 2007, *21*, 153–174.
- Marks, Melanie B. and Rachel T.A. Croson**, “Alternative rebate rules in the provision of a threshold public good: An experimental investigation,” *Journal of Public Economics*, 1998, *67*, 198–220.
- and –, “The effect of incomplete information in a threshold public goods experiment,” *Public Choice*, 1999, *99*, 103–118.
- Marwell, G. and R. Ames**, “Experiments on the provision of public goods II: Provision points, stakes, experience, and the free-rider problem,” *American Journal of Sociology*, 1980, *85* (4), 926–37.
- McFadden, Daniel**, “Conditional logit analysis of qualitative choice behavior,” in Paul Zarembka, ed., *Frontiers in Econometrics*, Academic Press, 1974, pp. 105–142.
- McKelvey, Richard D. and Thomas R. Palfrey**, “Quantal Response Equilibrium and Normal Form Games,” *Games and Economic Behavior*, 1995, *10*, 6–38.
- and –, “Quantal Response Equilibria for Extensive Form Games,” *Experimental Economics*, 1998, *1*, 9–41.
- , **Andrew M. McLennan, and Theodore L. Turocy**, “Gambit: Software Tools for Game Theory,” 2013. <http://gambit.sourceforge.net>.
- Mistiaen, J.A. and I.E. Strand**, “Location choice of commercial fishermen with heterogeneous risk preferences,” *American Journal of Agricultural Economics*, 2000, *82*, 1184–1190.
- Morris, Stephen and Hyun Song Shin**, “Global games, theory and applications,” *Working Paper*, 2001. Cowles Foundation for Research in Economics, Yale University.
- Munro, G.R. and A.D. Scott**, “The economics of fisheries management,” in A.V. Kneese and J.L. Sweeney, eds., *Handbook of Natural Resource and Energy Economics*, Elsevier Science Publishers, 1985.

BIBLIOGRAPHY

- Offerman, Theo and Joep adn Schram Arthur Sonnemans**, “Value orientations, expectations and voluntary contributions in public goods,” *The Economic Journal*, 1996, *106*, 817–845.
- , **Arthur Schram, and Joep Sonnemans**, “Quantal response models in step-level public good games,” *European Journal of Political Economy*, 1998, *14*, 89–100.
- Opaluch, J.J. and N.E. Bockstael**, “Behavioral modeling of fisheries management,” *Marine Resource Economics*, 1984, *1*, 105–115.
- Palfrey, Thomas R. and Jeffrey E. Prisbrey**, “Anomalous behavior in public goods experiments: How much and why?,” *The American Economic Review*, 1997, *87* (5), 829–846.
- Pauly, D.**, “Major trends in small-scale marine fisheries, with emphasis on developing countries, and some implications for the social sciences,” *Maritime Studies*, 2006, *4*, 7–22.
- Rapoport, Amnon**, “Individual Strategies in a Market Entry Game,” *Group Decision and Negotiation*, 1995, *4*, 117–133.
- , **Darryl A. Seale, and Lisa Ordonez**, “Tacit Coordination in Choice Between Certain Outcomes in Endogenously Determined Lotteries,” *The Journal of Risk and Uncertainty*, 2002, *25* (1), 21–45.
- , – , **Erev Ido, and James A. Sundali**, “Equilibrium Play in Large Group Market Entry Games,” *Management Science*, 1998, *44* (1), 119–141.
- Rogers, Brian W., Thomas R. Palfrey, and Colin F. Camerer**, “Heterogeneous quantal response equilibrium and cognitive hierarchies,” *Journal of Economic Theory*, 2009, *144*, 1440–1467.
- Sadovy, Y.**, “Trouble on the reef: The imperative for managing vulnerable and valuable fisheries,” *Fish and Fisheries*, 2006, *6*, 167–185.
- Sanchirico, J.N., U. Malvadkar, A. Hastings, and J.E. Wilen**, “When are no-take zones and economically optimal fishery management strategy?,” *Ecological Applications*, 2006, *16*, 1643–1659.
- Selten, Reinhard, Thorsten Chmura, Thomas Pitz, Sebastian Kube, and Michael Schreckenberg**, “Commuters route choice behavior,” *Games and Economic Behavior*, 2007, *58*, 394–406.

BIBLIOGRAPHY

- Sheremeta, Roman M.**, “Contest design: An experimental investigation,” *Economic Inquiry*, 2011, *49*, 573–590.
- Smith, M.D. and J.E. Wilen**, “Heterogeneous and correlated risk preferences in commercial fishermen: *The Perfect Storm* dilemma,” *The Journal of Risk and Uncertainty*, 2005, *31*, 53–71.
- Startz, R. and K.P. Tsang**, “Nonexponential discounting: A direct test and perhaps a new puzzle,” *The B.E. Journal of Macroeconomics*, 2012, *12*, 1–35.
- Sumaila, U.R.**, “Intergenerational cost-benefit analysis and marine ecosystem restoration,” *Fish and Fisheries*, 2004, *5*, 329–343.
- Turocy, Theodore L.**, “A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence,” *Games and Economic Behavior*, 2005, *51*, 243–263.
- , “Computing sequential equilibria using agent quantal response equilibria,” *Economic Theory*, 2010, *42*, 255–269.
- van Huyck, John B. and Ajalavat Viriyavipart**, “Testing risk dominance and payoff dominance in repeated global stag hunt games,” *Working Paper*, 2013. Department of Economics, Texas A&M University.
- , **Raymond C. Battalio, and Richard O. Beil**, “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure,” *The American Economic Review*, 1990, *80* (1), 234–248.
- Wang, M., M.O. Rieger, and T. Hens**, “An international survey on time discounting,” *Working Paper*, 2009.
- Weesie, Jeroen**, “Incomplete Information and Timing in the Volunteer’s Dilemma: A Comparison of Four Models,” *The Journal of Conflict Resolution*, 1994, *38* (3), 557–585.
- White, C. and C. Costello**, “Matching spatial property rights fishers with scales of fish dispersal,” *Ecological Application*, 2011, *21*, 350–362.

Appendices

Appendix A

Appendix to Chapter 1

A.1 Symmetric Bayes-Nash Equilibria

The threshold game has highly discontinuous payoffs, which makes it intractable on pencil and paper. This discontinuity is not incidental, it is a defining feature of this environment. Below the threshold nothing happens. Above the threshold, contributions are collected and benefits distributed in accordance with the payoff structure. Assuming the threshold is reached, any purely self-interested individual would hope to pledge as little as possible. Therefore, best response under complete information occurs at the discontinuity in payoffs generated by the threshold. Under incomplete information and with a discrete state space, there may be several discontinuities in payoffs. This is because additional contributions may lead to additional states of nature where the threshold is met.

Figure A.1: Generic Payoff Functions





Figure A.1 plots a generic payoff function for both information treatments above. As can be seen from both graphs, the best response occurs at a point of discontinuity, where no first order condition holds. In order to calculate the symmetric, pure-strategy Bayes-Nash equilibria, I perform a simple grid search across all potential contributions. I calculate the expected payoff function for a generic low type and high type player, assuming the other $n-1$ players choose (c_l, c_h) . If the optimal contribution by the generic high type and low type players both match the contributions of the other $n-1$ players, then I record that ordered pair as an equilibrium, otherwise the results are discarded. Using the endowments from the experiment, $w_i = 40$, the grid search is for discrete values over $(c_l, c_h) \in [0, 40] \times [0, 40]$. The results of this algorithm are available in Table 1.1 and plotted in Figure A.2 below.

Figure A.2: Sets of symmetric Bayes-Nash equilibria



To gain a deeper understanding of how this algorithm works, I will work out a few numerical examples. I will use the 4-player game for simplicity. Suppose I were to guess the solution

$$\{(c_l, c_h) | 2c_l + 2c_h = 80 \ \& \ 13 \leq c_l \leq 18.33\} \quad (\text{A.1})$$

I will illustrate by verifying a point within this set, and rejecting a point on the line that violates the restriction. Recall that each player's payoff function will have discontinuities at pivotal pledge levels where the public good is produced in an additional state of nature. Expected payoff maximization requires comparing the expected payoffs at these various points of discontinuity.

For example, suppose a generic player i is a low type, and all other players choose the strategy (15, 25). Below is Table A.1 listing the contribution levels required by player i to achieve the threshold in an additional state of nature, along

with the expected payoffs. It is abundantly clear that player i 's best response is to due her duty as a low type and choose $c_i^* = 15$.

Table A.1: Comparing Expected Payoffs at Discontinuities

Contribution	State	State Payoff	Expected Payoff
5	0Low/3High	15	$\frac{1}{8} \left[20 - 80 \left(\frac{5}{5+75} \right) \right] = 1.875$
15	1Low/2High	5	$\frac{1}{8} \left[20 - 80 \left(\frac{15}{15+75} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{15}{15+65} \right) \right] = 2.708$
25	2Low/1High	-5	$\frac{1}{8} \left[20 - 80 \left(\frac{25}{25+75} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{25}{25+65} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{25}{25+55} \right) \right] = -2.708$
35	3Low/0High	-15	$\frac{1}{8} \left[20 - 80 \left(\frac{35}{35+75} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{35}{35+65} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{35}{35+55} \right) \right] + \frac{1}{8} \left[20 - 80 \left(\frac{35}{35+45} \right) \right] = -9.723$

The same will be true at every point in the set of equilibria. However, it is important to note the second restriction in the guess above: $13 \leq c_l \leq 18.33$. Suppose I attempt to verify that $(10, 30)$ is not a Bayes-Nash equilibrium. I will not consider the case of a generic high type, so the restriction may be equivalently stated as $21.67 \leq c_h \leq 27$. Then the contribution levels that result in payoff

discontinuities will change as well, as shown in Table A.2. It appears that a generic high type would prefer to defect by only pledging 10.

Table A.2: Comparing Expected Payoffs at Discontinuities

Contribution	State	State Payoff	Expected Payoff
0	0Low/3High	20	$\frac{1}{8} \left[20 - 80 \left(\frac{0}{0+90} \right) \right] = 5.000$
10	1Low/2High	10	$\frac{1}{8} \left[20 - 80 \left(\frac{10}{10+90} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{10}{10+70} \right) \right] = 15.250$
30	2Low/1High	-10	$\frac{1}{8} \left[20 - 80 \left(\frac{30}{30+90} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{30}{30+70} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{30}{30+50} \right) \right] = 12.250$
50	3Low/0High	-30	$\frac{1}{8} \left[20 - 80 \left(\frac{50}{50+90} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{50}{50+70} \right) \right] + \frac{3}{8} \left[20 - 80 \left(\frac{50}{50+50} \right) \right] + \frac{1}{8} \left[20 - 80 \left(\frac{50}{50+30} \right) \right] = 2.679$

Identifying the exact location of these boundaries requires solving systems of high order polynomials. In practice, this is only practical using numerical methods, and the algorithms required are more computationally intensive than a the simple guess-and-check grid search that I perform. I parse the space $[0, 40] \times [0, 40]$ into steps of hundredths or thousandths. I then calculate the expected payoff that any generic player i would experience at any given strategy in $[0, 40]$

(also parsed into hundredths or thousandths) for both low and high payoff types. Finally, I record the intersection of the sets of strategies (c_l, c_h) such that a generic player i 's best response is to do as other low types when low and vice versa when high. This exhaustive search reveals the sets of equilibria to any reasonable level of numerical accuracy.

One apparent feature of these equilibria is that the maximum provision rate is decreasing in group-size. For example, groups of four can achieve an average provision rate of 86.75% while groups of eight could only achieve an average provision rate of 36.66%. Ramping up the algorithm to twenty and fifty players, I find the effect is quite dramatic. For twenty players, the maximum average provision rate is 13.16%, which is associated with seven or fewer low types. Meanwhile, groups of fifty may only achieve an average provision rate of 0.08%, or no more than sixteen low types. The dramatic drop in probabilities reflects the law of large numbers as group-size grows.

Appendix B

Appendix to Chapter 2

B.1 Theorems and Proofs

Theorem 1. *Assume that n_{in} and n_{out} players deterministically choose to enter or stay out of the market, respectively. Then there is a unique mixed-strategy equilibrium where all remaining players choose to enter with the same probability:*

$$p = \frac{c - n_{in} - 1}{n - n_{in} - n_{out} - 1} \quad (\text{B.1})$$

Proof. Define the remaining total number players, excluding an arbitrary player i , as $n_j = n - n_{in} - n_{out} - 1$, where each player chooses to enter with probability p_j , and define the total number of other mixing entrants as $S_j = \sum_{j=1}^{n_j} s_j$. Note that player i will be willing to mix if and only if the average net gain to entry is zero:

$$\sum_{S_j=1}^{n_j} \sum_{j=1}^{n_j} p_j^{S_j} (1 - p_j)^{n_j - S_j} \left[a + v_i (c - n_{in} - S_j - 1) \right] - r = 0 \quad (\text{B.2})$$

Imposing the restriction ($a = r$) used in this experiment provides the preference homogeneity that neutralizes any payoff magnitude effects:

$$\sum_{S_j=1}^{n_j} \sum_{j=1}^{n_j} p_j^{S_j} (1 - p_j)^{n_j - S_j} (c - n_{in} - S_j - 1) = 0 \quad (\text{B.3})$$

Notice that the term inside the parentheses may be separated as follows:

$$(c - n_{in} - 1) \sum_{S_j=1}^{n_j} \sum_{j=1}^{n_j} p_j^{S_j} (1 - p_j)^{n_j - S_j} - \sum_{S_j=1}^{n_j} \sum_{j=1}^{n_j} p_j^{S_j} (1 - p_j)^{n_j - S_j} S_j = 0 \quad (\text{B.4})$$

Within the first term, all of the summation collapses to unity, while in the second term, the summation collapse into a standard expected value calculation:

$$(c - n_{in} - 1) - \sum_{j=1}^{n_j} p_j = 0 \quad (\text{B.5})$$

Moving the second term to the right-hand-side yields the following equation:

$$c - n_{in} - 1 = \sum_{j=1}^{n_j} p_j \quad (\text{B.6})$$

Since the left-hand-side is a constant, and the right-hand-side depends (via exclusion) on player i , this statement will hold for each player i if and only if $p_i = p_j \forall i, j$:

$$n_j p_j = c - n_{in} - 1 \quad (\text{B.7})$$

Dividing both sides by n_j , and replacing this term with its definition $n_j = n - n_{in} - n_{out} - 1$, yields the familiar result from Rapoport et al. (1998):

$$p = \frac{c - n_{in} - 1}{n - n_{in} - n_{out} - 1} \quad (\text{B.8})$$

This is all quite trivial, since affine transformations of payoffs represent the same preferences according to von-Neumann Morgenstern utility: Define preferences by:

$$\tilde{u}_i(s_i, s_{-i}) = \begin{cases} (c - S) & \text{if } s_i = 1 \\ 0 & \text{if } s_i = 0 \end{cases} \quad (\text{B.9})$$

In this experiment, each player's payoff function may be written as the following affine transformation of these preferences

$$u_i = r + v_i \tilde{u}_i \quad (\text{B.10})$$

As such, the set of Nash equilibria is identical to the homogeneous case solved in Rapoport et al. (1998).

□

Theorem 2. *Assume that n_{in} and n_{out} players deterministically choose to enter or stay out of the market, respectively, regardless of payoff type. Define the remaining total number players, excluding an arbitrary player i , as $n_j = n - n_{in} - n_{out} - 1$, where each player j with payoff type k chooses to enter with probability p_j^k , and define the total number of mixing entrants in state k as $S_j^k = \sum_{j=1}^{n_j} s_j^k$. Then player i will be ex-ante indifferent between entry and staying out if and only if expected*

payoffs match those found with Nash equilibrium under complete information. In other words, the set of Bayes-Nash equilibria under private information is composed of the set of Nash equilibria under complete information, as well as any heterogeneous variation of these equilibria that is robust to cursed beliefs.

Proof. To begin, note that player i will be willing to mix if and only if the average net gain to entry is zero:

$$\sum_{k=1}^K \sum_{S_j^k=1}^{n_j} \sum_{j=1}^{n_j} q^k \left(p_j^k\right)^{S_j^k} \left(1-p_j^k\right)^{n_j-S_j^k} v_i \left(c-n_{in}-S_j^k-1\right) = 0 \quad (\text{B.11})$$

Since the indexation of players is arbitrary, the above restriction must hold for every player i who is mixing. Notice that the term inside the parentheses may be separated as follows:

$$\begin{aligned} v_i (c - n_{in} - 1) \sum_{k=1}^K \sum_{S_j^k=1}^{n_j} \sum_{j=1}^{n_j} q^k \left(p_j^k\right)^{S_j^k} \left(1-p_j^k\right)^{n_j-S_j^k} \\ - v_i \sum_{k=1}^K \sum_{S_j^k=1}^{n_j} \sum_{j=1}^{n_j} q^k \left(p_j^k\right)^{S_j^k} \left(1-p_j^k\right)^{n_j-S_j^k} S_j^k = 0 \end{aligned} \quad (\text{B.12})$$

Within the first term, all of the summation collapses to unity, while in the second term, the summation collapse into a standard expected value calculation:

$$v_i (c - n_{in} - 1) - v_i \sum_{k=1}^K \sum_{j=1}^{n_j} q^k p_j^k = 0 \quad (\text{B.13})$$

Moving the second term to the right-hand-side, and canceling the v_i 's, yields the following equation:

$$c - n_{in} - 1 = \sum_{k=1}^K \sum_{j=1}^{n_j} q^k p_j^k \quad (\text{B.14})$$

Since the q^k 's are exogenous and the left-hand-side of this equation is constant, and since the equilibria under complete information extend to private information, expected payoffs are fixed to the levels under complete information, and heterogeneity in entry probabilities is only permissible if it does not affect expected payoffs. Since these expected payoffs are associated with equilibria where actions reveal nothing about payoff type, the extension to heterogeneity must also be robust cursed beliefs. \square

Theorem 3. *Assume payoff types are distributed Bernoulli, and that the probability of being a low type is denoted by q . Then the average entry rate implied by the set of pure-strategy Nash equilibria represent an upper-bound, i.e., the average entry rate under pure strategies is at least as large as the average entry rate under all mixed-strategy equilibria.*

Proof. Start with the inequality implied by the theorem's conclusion:

$$n_{in} + (n - n_{in} - n_{out}) \cdot \left[\frac{c - n_{in} - 1}{n - n_{in} - n_{out} - 1} \right] \leq c \quad (\text{B.15})$$

Note that the expression above may be re-arranged to the following:

$$\frac{c - n_{in} - 1}{n - n_{in} - n_{out} - 1} \leq \frac{c - n_{in}}{n - n_{in} - n_{out}} \quad (\text{B.16})$$

Cross-multiplying yields:

$$(n - n_{in} - n_{out})(c - n_{in} - 1) \leq (c - n_{in})(n - n_{in} - n_{out} - 1) \quad (\text{B.17})$$

Distributing terms gives:

$$(c - n_{in})(n - n_{in} - n_{out}) - (n - n_{in} - n_{out}) \leq (c - n_{in})(n - n_{in} - n_{out}) - (c - n_{in}) \quad (\text{B.18})$$

Canceling the first term on each side and dividing through yields the following restriction:

$$\frac{c - n_{in}}{n - n_{in} - n_{out}} \leq 1 \quad (\text{B.19})$$

This is trivially true, given that the mixing rate is a probability:

$$0 < \frac{c - n_{in} - 1}{n - n_{in} - n_{out} - 1} < \frac{c - n_{in}}{n - n_{in} - n_{out}} < 1 \quad (\text{B.20})$$

□

Theorem 4. *Assume payoff types are distributed Bernoulli, and that the probability of being a low type is denoted by q . Then the following condition must hold in any mixed-strategy Bayes-Nash equilibrium:*

$$\sum_{j=1}^{n_j} p_j^l = \alpha + \beta \sum_{j=1}^{n_j} p_j^h \quad (\text{B.21})$$

Proof. From Theorem 2, the following relation must hold for any discrete type space:

$$c - n_{in} - 1 = \sum_{k=1}^K \sum_{j=1}^{n_j} q^k p_j^k \quad (\text{B.22})$$

Therefore, with only two types the equation simplifies to:

$$c - n_{in} - 1 = q \sum_{j=1}^{n_j} p_j^l + (1 - q) \sum_{j=1}^{n_j} p_j^h \quad (\text{B.23})$$

Re-arranging terms, the result obtains:

$$\sum_{j=1}^{n_j} p_j^l = \frac{1}{q} \cdot (c - n_{in} - 1) - \frac{1 - q}{q} \sum_{j=1}^{n_j} p_j^h \quad (\text{B.24})$$

The coefficients are determined by the exogenous parameters of the game:

$$\alpha = \frac{1}{q} \cdot (c - n_{in} - 1) \quad \& \quad \beta = -\frac{1 - q}{q} \quad (\text{B.25})$$

Given the design of this experiment, the parameters may take on the following values:

$$\alpha \in \left\{ \underbrace{1, \frac{4}{3}, \frac{2}{3}}_{\text{Both}}, \underbrace{\frac{3}{2}, \frac{6}{5}, \frac{8}{5}}_{\text{Large}} \right\} \quad \& \quad \beta = 1 \quad (\text{B.26})$$

□

Appendix C

Appendix to Chapter 3

C.1 Time preference experiment

Instructions:

This last part of the survey will only take a few minutes. For completing these last few survey questions, you will receive up to fl. 50. All you have to do is choose how you would like to be paid: sooner or later. There are 21 questions; each is a choice between a sooner payment and a later payment. After you answer all of the questions, you will draw a chip from this bag. The chips are numbered from 1 to 21 corresponding to the numbers of the questions. You will be paid according to your choice in the question whose number is on the chip. For example, if you choose the chip with the number 7 on it, you will receive the amount of money you chose in response to question 7 on the date mentioned in that question. Payments will be available for pick up on Fridays between 1pm and 5pm. I will personally

be there to give you your payment. Would you like to participate in this part of the interview?

Questions:

This Friday vs. 2 weeks from this Friday?

1. Would you like to receive fl. 50 this coming Friday or fl. 50 2 weeks from this coming Friday?
2. Would you like to receive fl. 48 this coming Friday or fl. 50 2 weeks from this coming Friday?
3. Would you like to receive fl. 44 this coming Friday or fl. 50 2 weeks from this coming Friday?
4. Would you like to receive fl. 40 this coming Friday or fl. 50 2 weeks from this coming Friday?
5. Would you like to receive fl. 35 this coming Friday or fl. 50 2 weeks from this coming Friday?
6. Would you like to receive fl. 30 this coming Friday or fl. 50 2 weeks from this coming Friday?
7. Would you like to receive fl. 20 this coming Friday or fl. 50 2 weeks from this coming Friday?

This Friday vs. 4 weeks from this Friday?

8. Would you like to receive fl. 50 this coming Friday or fl. 50 2 weeks from this coming Friday?
9. Would you like to receive fl. 48 this coming Friday or fl. 50 2 weeks from this coming Friday?
10. Would you like to receive fl. 44 this coming Friday or fl. 50 2 weeks from this coming Friday?
11. Would you like to receive fl. 40 this coming Friday or fl. 50 2 weeks from this coming Friday?
12. Would you like to receive fl. 35 this coming Friday or fl. 50 2 weeks from this coming Friday?
13. Would you like to receive fl. 30 this coming Friday or fl. 50 2 weeks from this coming Friday?
14. Would you like to receive fl. 20 this coming Friday or fl. 50 2 weeks from this coming Friday?

Two weeks from Friday vs. 4 weeks from this Friday?

15. Would you like to receive fl. 50 this coming Friday or fl. 50 2 weeks from this coming Friday?
16. Would you like to receive fl. 48 this coming Friday or fl. 50 2 weeks from this coming Friday?
17. Would you like to receive fl. 44 this coming Friday or fl. 50 2 weeks from this coming Friday?

18. Would you like to receive fl. 40 this coming Friday or fl. 50 2 weeks from this coming Friday?
19. Would you like to receive fl. 35 this coming Friday or fl. 50 2 weeks from this coming Friday?
20. Would you like to receive fl. 30 this coming Friday or fl. 50 2 weeks from this coming Friday?
10. Would you like to receive fl. 20 this coming Friday or fl. 50 2 weeks from this coming Friday?

Do you have a bank account?

Do you have a credit card?

Do you have a friend who would loan you fl. 50?

How much interest would they charge you?

C.2 Survey questions about management of coral reef resources

Gear Questions:

1. Should there be a limit to how small a mesh size can be on a trap?
2. Should rectangular holes be required in the side of fish traps to let small fish escape?
3. Should there be a limit to how small the mesh size can be on a seine net?
4. Should there be a limit to how small the mesh size can be on a gill net?

Appendix C. Appendix to Chapter 3

5. Should fishing with traps be banned?
6. Should snorkel fishing be banned?
7. Should fishing with seine nets be banned?
8. Should gill net fishing on reefs be banned?
9. Should gill net fishing be banned everywhere around the island?

Reserve Questions:

1. Do you support no fishing areas?
2. Do you think there should be more and/or larger no fishing areas?
3. Do you think there should be areas temporarily closed to fishing?
4. Should fishing on the reefs be banned?
5. Do you support no diving areas?
6. Do you think there should be more and/or larger diving areas?
7. Do you think there should be areas temporarily closed to diving?

Miscellaneous Management Questions:

1. Should there be a limit on the number of fishermen allowed to fish?
2. Should there be a limit to the number of divers allowed to dive?
3. Should catching juvenile fish be banned?

4. Should uncontrolled anchoring be banned if the government provides more buoys?
5. Are there any species you think should no longer be fished?