2011

Problems of uncertainty, learning, and welfare measurement in resource and environmental economics

Keith Shannon Evans
Iowa State University

Problems of uncertainty, learning, and welfare measurement in resource and environmental economics

by

Keith Shannon Evans

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Major: Economics

Program of Study Committee:
Joseph A. Herriges, Co-Major Professor
Catherine L. Kling, Co-Major Professor
Quinn Weninger, Co-Major Professor
John Miranowski
John R. Schroeter

Iowa State University
Ames, Iowa

2011

DEDICATION

I would like to dedicate this dissertation to my wife Courtney. I can hardly find the words to express my gratitude for your unfailing love and support along this journey. Without your strength this work would not be possible. Thank you.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGEMENTS

First and foremost, I am indebted to my major professors Joe Herriges, Cathy Kling, and Quinn Weninger. Thank you for the countless hours of guidance in writing this dissertation, and whose advice and criticisms will always be welcome.

A special thanks to Dr. V. Kerry Smith and Dr. Nicolai Kuminoff for their advice during my pre-doctoral fellow at Arizona State University. Their willingness to share their time, experience, and data helped make completion of the third essay possible. I hope to continue this relationship with them long into the future.

Many friends have helped and encouraged me throughout this process. There are several that I would like to thank in particular. Thank you, Jerome Dumortier, Luisa Menapace, Greg Colson, Abhishek Somani, Mindy Mallory, Carola Grebitus, Subhra Bhattacharjee, and Tristan and Kate Brown.

Finally, I would like to offer my deepest gratitude to my family for their patience with me during graduate school. It has been a long journey that is finally complete. To my parents Ron and Pat, my siblings Craig and Jennifer, and especially to my children Jade and James, thank you.

CHAPTER 1. GENERAL INTRODUCTION

1.1 Overview

This dissertation is a collection of three essays that focus on problems involving uncertainty, learning, and welfare measurement in resource and environmental economics. Each essay focuses on a separate aspect of these problems. The first and second essays focus on the behavior of fishermen in a dynamic setting of choice where there is uncertainty and learning. These two essays extend our understanding of the behavior and management of commercial fishermen on several fronts; including the formation of information sharing cooperatives, costly investment in information regarding productive fishing locations, and the rationalization of the fishing fleet after the introduction of a tradable rights-based management program. The third essay investigates the accuracy of predictions from the pure-characteristics vertical sorting model. This is a structural econometric model with the capability of predicting both spatial and distributional general equilibrium welfare effects from large-scale changes in the provision of the public good. The following is a brief description of each of these essays and an overview of the research contained within this dissertation.

The first essay, "Information Sharing and Cooperative Search in Fisheries" studies equilibrium search and learning in a dynamic fishing game that is played by independent fishermen and by members of an information sharing cooperative.

Several fishery management problems, from bycatch to the design of marine protected areas, have been linked to problems of information acquisition and information sharing in fisheries. An information sharing cooperative provides a governance structure that has the potential to lessen many of these issues. However, the design of fishery policy to promote such fishing cooperatives must account for the behavioral response of fishermen under this alternate incentive structure.

To date, the literature has largely ignored the internal governance structure of the fishing cooperative and instead modeled it as a single decision-maker. As such, a formal analysis of the search patterns of fishermen, including risk preference, the ability to share information and coordinate search, and the inclusion of private objectives of member fishermen is important.

The results of this paper suggest that independent fishermen do not internalize the full value of information and do not replicate first-best search patterns. An information sharing cooperative, though an improvement on independent effort, faces a free-riding problem as each member prefers that the costly search for information be undertaken by others. Devising contracts that result in optimal investment in information may be particularly challenging in fisheries, due to the club good characteristic of information, its costly acquisition, and the common property nature of the fishery resource. The addition of risk aversion significantly alters the equilibrium pattern of search and the demand for information through the introduction of disutility associated with experiencing risk.

The second essay, “Fleet Rationalization Under Individual Transferable Quotas,” develops a dynamic transition model to analyze the fleet size, configuration, and quota trading prices that emerge during the transition from an initially over-capitalized fishery to the long run cost-efficient fleet structure.

Individual transferable quotas (ITQs) offer an effective tool to address over-capitalization and rent dissipation that otherwise plague traditional command and control management approaches. ITQs provide incentives to efficiently utilize available economies of scale and over time align fleet harvest capacity with target catch levels. The long run efficiency gains that accompany rationalization are well documented, whereas, the time required to reduce the fleet size, an understanding of which vessels exit and which remain, the pricing of ITQs, and the efficiency implications during the transition are not well understood. With an ongoing push for rights-based management in U.S. fisheries, understanding the economic fundamentals of the transition is increasingly relevant.

The simulation results highlight the implications of the dynamic transition model regarding delayed-exit strategies, speculation, and the generation of rents. The results suggest that beliefs play a strong role on the transition. Uncertainty by fishermen over their relative cost-efficiency translates into uncertainty over quota trading prices. As such, a component of the ITQ asset's price is speculation over future trading prices. Heterogeneous uncertainty and learning also provide insight into observed patterns of vessel exit, generating new insight into the slow transition observed in U.S. fisheries. These results should be of interest to policy makers. The near universal practice of allocating the initial endowment of quota based on historic catch promotes delayed-exit strategies on cost-inefficient vessels and thereby prolongs the transition period during which the full efficiency benefits of ITQ management are unrealized.

The final paper, "Model Validation of the Pure-Characteristics Vertical Sorting Model," is the first attempt at external model validation for this structural econometric model.

There are well documented concerns with using the hedonic property value model to measure the welfare effect from large-scale or discrete changes in the provision of public amenities. The pure-characteristics vertical sorting model has been proposed as an alternative that may avoid many of these issues. Preference estimates along with the added structure are sufficient to predict equilibria that may emerge from a non-marginal change. However, while these models are capable of predicting how people and markets will adjust, we have no evidence on the accuracy of their predictions.

Using the cleanup of Luke Air Force Base (a deleted Superfund site in Maricopa County, AZ) as a quasi-experiment, data surrounding cleanup of this site is used to compare the observed and predicted sorting of households, prices, housing expenditures, and income. In-sample tests are also performed to see how well the data used in the estimation of model parameters fits the structure of the model. Welfare analysis provides an additional source of model validation, and highlights a proposed strength of the sorting model, namely its ability to analyze the distributional and spatial effects of government policy.

1.1.1 Dissertation Organization

The overall structure of the dissertation is ordered as follows. The next three chapters present the essays described above. Each of these chapters is considered a stand-alone paper. The dissertation is completed with a general discussion of the results from this body of work. All tables and figures are included within the appendices.

CHAPTER 2. INFORMATION SHARING AND COOPERATIVE SEARCH IN FISHERIES

A modified version of the paper to be submitted to a peer-reviewed journal

2.1 Introduction

Fishermen do not freely share valuable information. Studies conducted by anthropologists describe instead a culture of secrecy and deceit with regard to the location of productive fishing sites (Palmer (1990), Gatewood (1987, 1984), Andersen (1980), Orbach (1977), and Andersen and Wadel (1972)). Economists, on the other hand, have suggested that sharing information about good fishing sites will raise fishing profits by avoiding redundant search and avoiding congestion costs (Costello and Polasky (2008), Knapp (2008), Wilen and Richardson (2008), Costello and Deacon (2007), Lynham (2006), Gaspart and Seki (2003), Knapp and Hill (2003), and Wilson (1990)). Carpenter and Seki (2005) in fact show that Japanese shrimp fishermen who share information and profits realize higher harvest rates than non-cooperative shrimp fishermen.[1] These two strands of literature present a puzzling question: if cooperation and coordination among fishermen can yield economic benefits, why is information sharing and coordination of fishing activity rare in the fishery?

This puzzle is of added importance given renewed interest in promoting self-governance and property-rights based management in U.S. fisheries. The National Oceanic and Atmospheric Administration (2010) issued a Catch Share Policy in 2010 calling for the "adoption of catch shares where appropriate in fishery management and ecosystem plans and their amendments, and will support the design, implementation, and monitoring of catch share programs."

[1] The potential gains of cooperation by fishermen have been recognized in U.S. fisheries management legislation. The Fisheries Collective Marketing Act (FCMA) of 1934 and the American Fisheries Act of 1998 allow the formation of fishing cooperatives. For a deeper discussion see Kitts and Edwards (2003).

NOAA's policy specifically includes provisions to assign portions of the total allowable catch to fishing cooperatives. Fishery management councils have responded to this policy. The North Pacific Fishery Management Council (2011) requested a discussion paper on the use of fishing cooperatives in the Gulf of Alaska Pollock fishery to combat problems of incidental catch of Chinook Salmon. As outlined in the discussion paper, the formation of fishing cooperatives is "intended to facilitate *information sharing and fleet coordination* that could be important to achieving Chinook avoidance" (pp. 3, emphasis added). In a similar effort, amendment 16 to the Northeast Multispecies Fisheries Management Plan, developed by the New England Fishery Management Council, pushes for expansion of the use of fishing cooperatives (New England Fishery Management Council (2010)). Under amendment 16, the National Marine Fisheries Service is responsible for determining the number of fishing cooperatives that conduct harvesting operations and the annual total harvests of a host of groundfish species. These examples highlight the importance of examining the efficiency of cooperative efforts among fishermen.

Several management problems, from the bycatch of unwanted species to spatial closures to the design of marine protected areas, have been linked to problems of information acquisition and information sharing in fisheries (Abbott and Wilen (2010), Haynie et al. (2009), Marcoul and Weninger (2008), Costello and Deacon (2007), Gilman et al. (2006), and Curtis and McConnell (2004)). An information sharing cooperative provides a governance structure that has the potential to lessen many of these issues. However, to design effective management policy, it is vital to understand the behavioral response of fishermen. As such, we examine active search with learning, including the consequences of risk preference, in a setting of multiple fishermen. In this paper we address four main research questions: (1) what is the first-best search pattern, (2) under what conditions do independent fishermen replicate the first-best, (3) do information sharing and effort coordination cooperatives necessarily replicate the first-best, and (4) how does risk preference alter the search problem?

To answer these questions, we develop a simple dynamic fishing game and examine the incentives to undertake costly search and share information across different levels of risk preference. We compare equilibrium levels of exploration, learning, congestion, and harvesting performance under various institutional structures ranging from independent fishermen to a stylized fishing cooperative that utilizes simple contracts to coordinate the activities of member fishermen. We are able to investigate the importance of the internal governance structure within the cooperative for realizing gains from information sharing and coordination of fishing activities.

Our results show that the conditions, e.g., congestion costs and fishermen's beliefs about uncertain payoffs at competing fishing sites, under which benefits from information sharing and fishing coordination arise are special. Not surprisingly, cooperation is most valuable when learning is relatively complete, when congestion penalties are large, when information transmission among fishermen is costless, and when information about the true location of productive fishing sites does not quickly decay. A less obvious finding is the importance of the free-riding problem. Our model emphasizes the role of active and costly information acquisition by fishermen. Steaming to an uncertain fishing site to learn about its true productivity utilizes fuel, labor and bait, and importantly, implies time lost fishing at some other site. Once acquired, information about productive fishing sites is an excludable and non-rival good, or club good. To prevent free-riding, an information sharing group must distribute informational rents among its members, while maintaining incentives to undertake costly search for fish.[2] We show how the extent of the free-riding problem depends on the contract used to divide information acquisition costs and information rents among cooperative members. Devising contracts that result in optimal investment in information may be particular challenging in fisheries, due to the club good characteristic of information, its costly acquisition, and the common property nature of the fishery resource.[3] Our results provide an explanation for the paucity of information sharing and coordinating efforts in fisheries, and highlight the importance of risk preference in predicting the response of fishing behavior to regulation. This paper provides baseline results, and a framework for further research into the efficiency of alternative management tools.

[2]Buchanan (1965) addresses the efficiency conditions for club goods, arguing the formation of a club, with a membership fee or cost sharing, can realign incentives to prevent free-riding.

[3]According to U.S. Department of Agriculture (2009), there were 37 fishing cooperatives operating in the U.S. in 2009. This was a drop from 70 active fishing cooperatives operating in 1980; a meager 5% of fishers at the time; see Kitts and Edwards (2003) and Garland and Brown (1985).

2.2 Background and Previous Literature

Information sharing in fisheries has been studied extensively by anthropologists and to a lesser extent by economists. A variety of theories have emerged to explain why fishermen may or may not share information. Wilson (1990) suggests fishermen discover productive and profitable fishing sites by chance and must then decide whether or not to divulge the location of the site to other members of a fishing *club*. He suggests that information will be shared if the cost to the fisherman who discovers the productive fishing site is not "too large." Wilson argues information sharing costs will be small when the number of club members is small, when club members have high catching capacity relative to the size of the stock at the site, and when transaction costs accompanying information transmission are small. Wilson suggests further that the benefits of joining the information sharing club increase when club members are equally skilled at finding fish and therefore are likely to reciprocate, i.e., share valuable information, at some future date.

The anthropology literature offers similar explanations for information sharing among fishermen. Gatewood (1984), Orbach (1977), Stiles (1972), and others suggest that fishermen will be more secretive, i.e., less inclined to share information, when the direct cost of disclosing information is high. Palmer (1990) finds some supporting evidence for this hypothesis among Maine lobstermen. His study of radio transmissions finds that lobster fishermen are more inclined to discuss detailed information about lobster size and fishing locations during pre-molting periods when lobster are hidden in rocks and less accessible to trap gear. In other words, information is shared more freely when it has little value.

Despite the prevalence of secrecy, the economics and anthropology literatures agree that information sharing and coordination should be profitable. Costello and Deacon (2007) suggest that joining an information sharing group can prevent redundant search that will otherwise occur when fishermen act independently (see also Wilson (1990)). A coalition of fishermen, such as a fishing cooperative, may be able to economize on search through specialization; with one or more members carrying out search activities while others specialize in harvesting fish. Independent fishermen on the other hand must forgo such specialization and search out productive fishing locations for themselves.[4] Another motive for coordinating fishing activities is to avoid congestion costs that arise when multiple fishermen happen to visit the same fishing site.[5]

A common theme in previous literature is the criterion of quid pro quo information sharing; fishermen will voluntarily disclose valuable information if and only if they receive something in return. But the criteria for forming an information sharing group must be more stringent. If there is to be a net expected gain from joining, it must be the case that expected profits per club member exceed the expected profits when fishermen act alone, i.e., there must be a return to scale in information gathering, or alternatively, an information synergy.[6]

In sum, past research on information sharing and effort coordination in fisheries has yielded insights that are sensitive to three underlying assumptions. First, researchers have assumed that fishermen are myopic, passive learners (Costello and Deacon (2007), Lynham (2006), and Smith (2000)). This behavioral assumption precludes costly investment in information and overlooks the free-riding problem.[7] Second, researchers have ignored the internal governance structure of a fishing cooperative, and have assumed that the incentives of individual cooperative members are fully aligned. Without this potential source of friction, the co-operative trivially coincides with the first-best outcome in terms of information acquisition and net returns to fishing. Third, models of active search and learning have focused on the case of risk neutral preferences (Marcoul and Weninger (2008) and Mangel and Clark (1986, 1983)), which forces the value of

[4] A few articles suggest that fishing cooperatives present a first-best solution for fishery management, see Deacon et al. (2008), Matulich et al. (2001), and Criddle and Macinko (2000).

[5] Congestion can lower profits when boats interfere with each others fishing activities. Alternatively, there can be localized temporary stock effects that reduce fishing profits. Alaskan halibut fishermen occasionally arrive at a site that has been fished a day or two earlier. They will learn this only after they pull their gear and find that fishing is exceptionally poor. The stock effect is temporary, and after a few days of rest, halibut from surrounding areas will redistribute back into the site. This temporary stock effect is consistent with the behavior MacCall (1990) predicts in his basin model. Coordinating the location of fishing activities could avoid the costs of congesting a site or fishing it before it has recovered from a previous visit.

[6] Gatewood (1984) identifies a synergy in information sharing in the Alaskan purse seine fishery. Managers control fleetwide harvest using periodic openings and closures the fishery. An impending opening is announced usually a day or two in advance. Gatewood finds that seine fishermen form information sharing groups noting that "While it is true that one boat can scout as wide an area in four days as four boats can in one day, the utility of the information collected by the four boats scouting the day before the opening is much greater, provided they share what each has observed" (pp. 362). It is important to note that the information gathering setting for salmon seiner's differs from the one that is modeled in this paper.

[7] Active, forward-looking learners recognize the connection between the information gathered currently and the quality of the choices they can make in the future. Marcoul and Weninger (2008) and Mangel and Clark (1986, 1983) demonstrate that under a repeated choice setting, a fisherman that is an active learner may choose to fish a site with a lower expected payoff to learn information about the site.

the information to grow monotonically with uncertainty. Other specifications of risk preference introduce non-monotonicities that generate alternative search patterns.

The model that we present next relaxes all three assumptions. We consider explicitly the cost of gathering information and synergies that arise from the public good nature of information. We then extend our analysis to consider fishing behavior in the case where the profits of a fishing cooperative are distributed among its members using simple wage contracts. Our framework allows us to derive insights for the behavior of individual cooperative members, and overall performance of fishing cooperatives that are absent from earlier literature.

2.3 Model

Our model features four key elements of the information sharing problem in fisheries: (1) uncertainty over true payoffs at competing fishing locations, (2) the opportunity to actively learn and reduce uncertainty over time, (3) gains from shared information, and (4) gains from coordinating site choices to avoid congestion.

We introduce a two-period game played by two risk neutral fishermen, hereafter referred to as the Row and the Column player. The case of risk aversion is considered later. To ease notation, a breve will be placed above variables associated with the Column player, a prime is used to distinguish second period values, and a subscript j differentiates the two fishing sites.

The timing of the game is illustrated in Figure A.1. Site choices are made at dates $t = 0$ and $t = 1$. Only one site can be fished per period. We assume that fishing at one of the sites is always preferred to not fishing at all.

At the time sites are chosen, fishermen are uncertain about fishing success and payoffs. Uncertainty derives from unpredictable variation in stock abundance and production conditions; e.g. the movement of fish within a site, and/or random weather and tides. Payoff uncertainty is modeled as random beliefs about true payoffs. We assume that beliefs follow a normal distribution. At $t = 0$, the Row player's beliefs about the true payoffs are summarized as, $b = \{\mu_1, \nu_1, \mu_2, \nu_2\}$ where μ_1 denotes the mean payoff at site 1, ν_1 denotes the variance of the mean payoff at site 1, and so on. Following our notational convention, the date $t = 0$ beliefs of the Column player are denoted as $\breve{b} = \{\breve{\mu}_1, \breve{\nu}_1, \breve{\mu}_2, \breve{\nu}_2\}$.

Following the site choice, fishing yields a payoff realization,

$$s_j = u_j + \epsilon_j$$

where u_j is the unobservable true payoff and ϵ_j is a random term. We assume the distribution of ϵ_j is known by the fishermen; $\epsilon_j \sim N(0, \nu_s)$, where $\nu_s > 0$ denotes the variance of payoff noise. Signal noise and therefore payoffs are independently distributed across sites.

The payoff realized in the first fishing period is used to update beliefs about the true payoff distribution following Bayes' rule. Our assumptions on the distribution of beliefs and payoffs imply that updated beliefs follow a normal distribution; see Gelman et al. (2003) and Appendix A.1 for additional details.

A player's action is their choice of fishing site in each period. We allow for mixed strategies and denote the probability the Row player selects site j in period 1 by $a_j \in [0, 1]$. The date $t = 0$ action of the Row player is $a = (a_1, a_2)$ where $a_1 + a_2 = 1$. Similarly the Column players action is $\breve{a} = (\breve{a}_1, \breve{a}_2)$ with $\breve{a}_1 + \breve{a}_2 = 1$.

Second period payoffs are discounted by a factor $\delta \in [0, 1]$, which we assume is common across fishermen. We will focus on the case where $\delta = 1$, no discounting of future payoffs, and $\delta = 0$, the case of myopic fishermen. Our results are not affected qualitatively under the more realistic case where $\delta \in (0, 1)$.

Lastly, a congestion penalty, denoted $\kappa \geq 0$, is deducted from the realized fishing payoff in the event that both fishermen select the same site in a fishing period.[8] The congestion penalty is fixed and common across fishermen and sites. In the analysis that follows we often set $\kappa = 0$ to isolate the role of information sharing in the model. When $\kappa > 0$ the payoffs for each fisherman depend on the action of his counterpart.

All site choices are based on *beliefs* about true fishing payoffs and, for the case of non-cooperating fishermen, beliefs about the strategy of the other fisherman. For example, consider the date $t = 0$ expected payoff for the non-cooperating Row player. Selecting site j yields an expected lifetime payoff that includes a first period expected payoff from fishing site j, plus the discounted expected payoff from responding optimally in the second fishing period. The

[8]Congestion penalties are common in the literature on common pool resource exploitation; see Smith (1968).

first period expected payoff is increasing in the Row fisherman's belief about the true payoff at site j, and non-increasing in his conjecture about the likelihood that the Column player will also fish the site. This conjecture depends on his *beliefs about the beliefs* held by the Column player, which is a standard feature of Bayesian games; see Gibbons (1992). The discounted expected payoff in the second fishing period depends on updated beliefs, and correspondingly the information gathered in the first fishing period, as well as the Row fisherman's conjecture about the updated strategy of his rival. Fishermen understand that fishing a site provides information that can guide future site choices. Part of the benefit of visiting a site is therefore the value of the information contained in the payoff signal.

Much of the analysis that follows will focus on decisions to undergo costly information acquisition. In the sequel we define an *investment in information* as follows.

Definition 1. *An* **investment in information** *is a site choice that yields a lower expected payoff than a competing site, but high information value. The cost of acquiring information at site j, rather than fishing site i, is the net expected payoff sacrificed in the current period by fishing at the site; e.g.* $\mu_i - \mu_j$.

We begin by characterizing the first-best outcome of the two-period search problem. Under the first-best, the two fishermen coordinate search efforts and freely share all information with a common goal of maximizing combined fishing profits. The first-best choices and payoffs provide a benchmark from which to compare outcomes under alternative governance structures, i.e., independent fishermen and our stylized fishing cooperative.

2.3.1 First-Best Coordination and Information Sharing

We imagine a situation where a manager, or sole owner, directs each fisherman to a particular site in each period, with the objective of maximizing total expected payoffs. Optimal site choices and payoffs are solved recursively.

At date $t = 1$, two payoff signals are available under three possibilities: both fishermen fished site 1, both fished site 2, and each fisherman fished a separate site in the first period. We identify the optimal policy and payoffs for the case where both fishermen fished site 1. The analysis of the remaining cases follows analogously and to conserve space is not repeated.

With one fishing period remaining, the manager uses his updated beliefs to compare the expected payoffs of the different combinations of site choices. As the manager is indifferent which fisherman is sent to a site, the optimal site choices take a simple form,

$$\begin{cases} (a_1' = 1, \breve{a}_1' = 1) & \text{if} \quad \mu_1' - 2\kappa > \mu_2' \\ (a_1' = 0, \breve{a}_1' = 0) & \text{if} \quad \mu_1' < \mu_2' - 2\kappa \\ (a_1' = 1, \breve{a}_1' = 0) & \qquad \text{otherwise} \end{cases} \tag{2.1}$$

Belief updating formulas are linear functions of s_1 and $\breve{s}_1$, therefore the policy in (2.1) is easily inverted. It will be convenient to express the date $t = 1$ policy in terms of realized signals. We divide the set of all site 1 signals into the following mutually exclusive subsets;

$$\begin{aligned} S_{11} = &\quad \{(s_1, \breve{s}_1) |\ \mu_1' - 2\kappa > \mu_2'\} \\ S_{22} = &\quad \{(s_1, \breve{s}_1) |\ \mu_1' < \mu_2' - 2\kappa\} \\ S_{12} = &\quad \{(s_1, \breve{s}_1) \notin S_{11} \cup S_{22}\} \end{aligned} \tag{2.2}$$

From the optimal policy in (2.1), total expected payoffs at date $t = 1$, expressed as a function of realized signals, is given as,

$$v(s_1, \breve{s}_1) = \begin{cases} 2\mu_1' - 2\kappa & \text{if } (s_1, \breve{s}_1) \in S_{11} \\ 2\mu_2' - 2\kappa & \text{if } (s_1, \breve{s}_1) \in S_{22} \\ \mu_1' + \mu_2' & \text{if } (s_1, \breve{s}_1) \in S_{12} \end{cases}$$

Note that $v(\cdot)$ is random at the time that sites are chosen.

To calculate the expected payoffs for the remaining two cases, i.e., where both fishermen are directed to site 2 or fishermen are directed to different sites, involves different updating formulas but otherwise similar calculations.

Next we step back to consider date $t = 0$ site choices. At $t = 0$, the manager has beliefs about the distribution of payoffs at each site and can calculate optimal site choices and second-period payoffs (equations (2.1) and (2.2)). Calculating two-period expected payoffs contingent on the manager's assignment of fishing locations is straightforward. For example, suppose the manager decides to send both fishermen to site 1 in the first fishing period. The two-period expected payoff from this action, denoted V_{11}, is equal to,

$$V_{11} = 2\mu_1 - 2\kappa + \delta \iint v(s_1, \breve{s}_1) d\Phi(\breve{s}_1) d\Phi(s_1), \tag{2.3}$$

where Φ denotes the normal cumulative density function.[9]

The remaining expected payoffs follow similarly. Sending both fishermen to site 2 in the first fishing period yields the two-period expected payoff,

$$V_{22} = 2\mu_2 - 2\kappa + \delta \iint v(s_2, \breve{s}_2) d\Phi(\breve{s}_2) d\Phi(s_2), \tag{2.4}$$

and a decision to send one fisherman to each site yields the expected payoff,

$$V_{12} = \mu_1 + \mu_2 + \delta \iint v(s_1, \breve{s}_2) d\Phi(\breve{s}_2) d\Phi(s_1). \tag{2.5}$$

The total expected payoff under the first-best is given as,

$$V^{FB} = \max \{V_{11}, V_{12}, V_{22}\}.$$

We note some key features of the optimal policy and payoffs. First, in the first-best case with full information and coordination there is no strategic interaction; site choices are deterministic. This will not be the case for independent fishermen considered next. Second, although suppressed for notational convenience, the site choice policy function and value function are determined solely by beliefs about true payoffs at competing sites.

The value functions provide insight into the first-best search pattern. Consider the case where $\mu_1 > \mu_2$. The sole owner will invest in a single signal from site 2 if and only if $V_{12} > V_{11}$ and $V_{12} > V_{22}$. This simplifies to bound investment costs where exploring site 2 is optimal.

$$\delta \Big\{ E[v(s_2, \breve{s}_2)] - E[v(s_1, \breve{s}_2)] \Big\} - \kappa < \mu_1 - \mu_2 < \delta \Big\{ E[v(s_1, \breve{s}_2)] - E[v(s_1, \breve{s}_1)] \Big\} + \kappa \tag{2.6}$$

That is, a single investment will occur when the investment cost if bound between the discounted total expected gain from the first and second signals. Then it follows that the manager will send both vessels to explore site 2 if and only if the added cost is less than the discounted total expected gain from the second signal.

Three insights follow directly from equation (2.6).

[9]For notational convenience, note that $\int \phi(x)dx$ can be rewritten as $\int d\Phi(x)$ where $\phi(x) = \frac{d\Phi(x)}{dx}$.

Result 1. (Myopic sole-owner and congestion effects).

1. *The sole owner recognizes the non-rival nature of information and incorporates its full cost and benefit into the investment decision.*

2. *A myopic sole owner* ($\delta = 0$) *ignores the value of information and simply compares current period expected payoffs at competing sites; e.g. no investment in information.*

3. *A positive congestion penalty* ($\kappa > 0$) *increases the region where a single investment in information occurs, while reducing the region for a second investment.*

2.3.2 Independent Fishermen

Independent fishermen choose sites simultaneously at date $t = \{0, 1\}$. We assume that at date $t = 1$ each knows only his counterpart's period 1 site choice; realized payoffs are private information. To proceed we must specify what each fisherman believes about the beliefs held by his counterpart. We assume fishermen share common date $t = 0$ beliefs about true payoffs.[10] We therefore solve for a symmetric Perfect Bayesian Nash Equilibrium assuming initial beliefs satisfy $b = \breve{b}$. The model is again solved recursively.

Consider the following date $t = 1$ strategy profile for the Row player,[11]

$$a'(s_j) = \begin{cases} \text{Fish site } j & \text{if} \quad s_j > s^U \\ (a'_1,\ a'_2) & \text{if} \quad s^L \leq s_j \leq s^U \\ \text{Fish site } i & \text{if} \quad s_j < s^L \end{cases} \tag{2.7}$$

We assume the Column player plays a similar strategy, though his signal thresholds, $\breve{s}^U$ and $\breve{s}^L$, are allowed to be different.

In a Bayesian Nash Equilibrium, each player's action must be optimal subject to their Bayesian updated belief about the strategy of their rival. Upon observing a private fishing signal in period 1, each player updates beliefs to form a best response to their rival's strategy.

[10] If this were not the case, each fishermen must believe that his counterpart holds biased beliefs. Moreover, neither fisherman should believe that his counterpart holds unbiased beliefs that differ from his own. Otherwise his own beliefs should be updated to incorporate the unbiased information. While our model can accommodate different initial beliefs, an explanation for why one fisherman believes his counterpart has beliefs different than his own is not obvious.

[11] There may be other strategy profiles. We choose to focus on the strategy presented in equation (2.7) as it is flexible to both pure and mixed strategies.

The thresholds s^U and s^L demarcate three regions. Suppose both fishermen fished site j in the first fishing period. If s_j is large, the Row player is encouraged to return to site j in the second fishing period. However, if the Column fisherman also fished site j, the Row player's private information also provides information on the probable private information of the Column player. A high s_j indicates a higher likelihood that the Column players may also return to site j in the second fishing period. The threshold s^U is the value of s_j that makes the Row player just indifferent between fishing site j as a pure strategy and mixing between the sites.

By similar arguments, a particularly low signal from site j will encourage the fisherman to switch to site i in the second fishing period. s^L is the threshold value for s_j that makes the Row player just indifferent between fishing site i as a pure strategy and mixing between the sites. Signals in the intermediate range, $s^L < s_j < s^U$, require a mixed strategy be played in the second fishing period. With an intermediate signal the Row player must deduce that it is likely the Column player also received a signal in the intermediate range. Updated beliefs suggest that congestion is likely and therefore a mixing strategy is preferred.

It should be obvious that the thresholds s^L and s^U will be functions of prior beliefs and the congestion penalty, as well as the first period site choice of the other fishermen. Derivation of the threshold values s^U, s^L, $\breve{s}^U$, and $\breve{s}^L$ is left to Appendix A.2.

Each player's strategy profile determines his $t = 1$ equilibrium site choices for any $t = 0$ fishing signal. For example, the Row player's expected payoff at date $t = 1$, conditioned on both fishermen fishing site 1 in the first period, is given as,

$$v(s_1|\breve{j} = 1) = \begin{cases} \mu_1' - \breve{a}_1'\kappa & \text{if} \quad s_1 > s^U \\ \mu_2' - \breve{a}_2'\kappa & \text{if} \quad s_1 < s^L \\ a_1'(\mu_1' - \breve{a}_1'\kappa) + a_2'(\mu_2' - \breve{a}_2'\kappa) & \text{if} \quad s^L \leq s_1 \leq s^U \end{cases}$$

where $\breve{j}$ denotes the site choice of the Column player.

We now step back and consider date $t = 0$ site choices. The strategy profile introduced above indicates a best response to any payoff realization. This is sufficient to determine the two-period conditional expected payoff; conditioned on both player's site choice. For example, following the notation of the previous section, the Row fisherman's expected payoff at site 1, conditional on the Column player also selecting site 1 is,

$$V_{11} = \mu_1 - \kappa + \delta \int v(s_1 | \breve{j} = 1) d\Phi(s_1) \tag{2.8}$$

The Row player conditions his expected future payoffs on his best belief about the signal observed by the Column player. That is, his observation of s_1 provides valuable information about the distribution that generated the Column player's private information, $\breve{s}_1$. This information is incorporated into his conjecture about the probable action of the Column player. The remaining cases are calculated similarly and take the following form,

$$V_{12} = \mu_1 + \delta \int v(s_1 | \breve{j} = 2) d\Phi(s_1) \tag{2.9}$$

$$V_{21} = \mu_2 + \delta \int v(s_2 | \breve{j} = 1) d\Phi(s_2) \tag{2.10}$$

$$V_{22} = \mu_2 - \kappa + \delta \int v(s_2 | \breve{j} = 2) d\Phi(s_2) \tag{2.11}$$

Using the conditional expected payoffs in equations (2.8) - (2.11), it is a simple exercise to construct a 2×2 normal form representation of the date $t = 0$ site choice game. The Perfect Bayesian Nash Equilibrium (PBNE) fishing strategy for the Row fisherman is the Nash Equilibrium of this normal form game. Then the date $t = 0$ strategy takes the form,

$$a_1 = (1 - a_2) = \begin{cases} 1 & \text{if } V_{11} > V_{21} \text{ and } V_{12} > V_{22} \\ 0 & \text{if } V_{22} > V_{12} \text{ and } V_{21} > V_{11} \\ \frac{V_{22} - V_{12}}{V_{11} - V_{21} + V_{22} - V_{12}} & \text{otherwise} \end{cases}$$

The maximum expected payoff for an independent fisherman is given as,

$$V^I = a_1 \Big(\breve{a}_1 V_{11} + \breve{a}_2 V_{12} \Big) + a_2 \Big(\breve{a}_1 V_{21} + \breve{a}_2 V_{22} \Big)$$

These value functions provide insight into the search pattern of independent fishermen and offer a comparison with the first-best. Consider the case where $\mu_1 > \mu_2$. An independent fisherman will invest in a signal from site 2 (as a pure strategy) if and only if the two-period expected value of fishing site 2 exceeds that of fishing site 1, or $\breve{a}_1 V_{21} + \breve{a}_2 V_{22} > \breve{a}_1 V_{11} + \breve{a}_2 V_{12}$.

Again, this simplifies to provide a bound on investment costs where exploring site 2 is optimal.

$$\begin{aligned}\mu_1 - \mu_2 < \delta\big[&\breve{a}_1\big\{E[v(s_2|\breve{j}=1)] - E[v(s_1|\breve{j}=1)]\big\}\\ &+ \breve{a}_2\big\{E[v(s_2|\breve{j}=2)] - E[v(s_1|\breve{j}=2)]\big\}\big]\\ &+ (\breve{a}_1 - \breve{a}_2)\kappa.\end{aligned} \tag{2.12}$$

The right-hand side of (2.12) represents the discounted expected net gains from exploring the less certain site. This inequality highlights the dependence of a fisherman's investment decision on his conjecture about his rival's fishing strategy.

By similar argument, an independent fisherman will play a mixed strategy when fishing a site is only optimal when it is not congested ($V_{21} > V_{11}$ but $V_{22} < V_{12}$). His mixed fishing strategy balances the potential gains of information with the penalty for congesting a fishing location, a coordination game. After a little algebra, independent fishermen play this coordination game when the investment cost is bound between,

$$\begin{aligned}\mu_1 - \mu_2 &< \delta\Big\{E[v(s_2|\breve{j}=1)] - E[v(s_1|\breve{j}=1)]\Big\} + \kappa\\ \mu_1 - \mu_2 &> \delta\Big\{E[v(s_2|\breve{j}=2)] - E[v(s_1|\breve{j}=2)]\Big\} - \kappa.\end{aligned} \tag{2.13}$$

Four results follow directly from equations (2.12) and (2.13).

Result 2. (Myopic independent fishermen and congestion effects).

1. *The independent fisherman only incorporates his private cost and benefit of information into the investment decision.*

2. *A myopic decision maker* ($\delta = 0$) *ignores the value of information and simply compares current period expected payoffs at competing sites; e.g. no investment in information.*

3. *A positive congestion penalty* ($\kappa > 0$) *increases the region where an independent fisherman plays a mixed strategy between exploiting site 1 and exploring site 2. This mixed strategy will depend on both the fisherman's beliefs about true payoffs at competing sites as well as his beliefs about the beliefs of the rival fisherman.*

4. *If the congestion penalty is zero* ($\kappa = 0$)*, the strategy profile collapses to only allow pure strategies, which will be independent of the actions of the rival fisherman.*

2.3.3 The Inefficiency of Independent Search

Under the first-best, the full costs and benefits of acquiring information are internalized. In contrast, independent fishermen do not share information once it is acquired, yet bear the full cost of collecting it. Comparing these investment criteria, it is evident there will exist initial beliefs for which investment in information at a particular site will be privately optimal for both independent fishermen, but the two privately observed signals are redundant relative to the first-best. In a like manner, beliefs will exist where an independent fisherman would choose not to gather costly information that is collected under the first-best since the information value is higher when it is shared by multiple fishermen.

Analytic solutions for equilibrium site choice strategies are not available and are calculated numerically. Figure A.2 reports these strategies over a range of date $t = 0$ beliefs.[12] The axes depict the ratio of the mean payoff at site 1 relative to the mean payoff at site 2, and the ratio of payoff uncertainty at site 2 relative to payoff uncertainty at site 1. At the origin payoff beliefs at the two sites are identical. Except at the origin, $\mu_1/\mu_2 > 1$ and $\nu_2/\nu_1 > 1$, i.e., expected payoffs are higher at site 1 but uncertainty and potential information value is higher at site 2. A choice to visit site 2 therefore represents a costly investment in information; depicted by $a_2 = 1$ and $\breve{a}_2 = 1$ in the figure. The figure reports the total number of visits to site 2 and thus compares aggregate investment in information by independents and under the first-best. Panel (a) and (c) in Figure A.2 depict the aggregate visits to site 2 under the first-best; panel (b) and (d) show visits to site 2 by independent fishermen. The top panels (a) and (b) illustrate the case of no congestion penalty and therefore isolate the role of information sharing. In panels (c) and (d) κ is set to 1% of μ_2, allowing strategic action by independent fishermen.

Consider the results with no congestion penalty. Marcoul and Weninger (2008) show that for a given investment cost their will exist a threshold level of mean payoff uncertainty at which an investment in information becomes optimal. The equilibrium policies in panels (a) and (b) identify this threshold value (the dark solid line in the figure), separating the belief space into a region where information investment does and does not occur. Comparing independent

[12]The simulation is parameterized as follows. Beliefs at the origin are fixed to be equal; $\mu_1 = \mu_2 = 100$ and $\nu_1 = \nu_2 = 3$. Information about the productivity of fishing locations is derived from noisy signals with constant variance $\nu_s = 2$. Alternative choices for simulation parameters did not qualitatively affect the results.

fishermen to the first-best, it is apparent that the threshold uncertainty value is smaller under the first-best policy. Intuitively, lower information value is required to offset investment costs since the information is shared.

Notice that when μ_1/μ_2 or ν_2/ν_1 becomes large, the actions of independent fishermen coincide with the first-best. If the cost of collecting information at site 2 becomes too large, or the value of information too great, the decision of where to fish is no longer controversial. However, when there is *tension* between the two sites, the actions of independent fishermen differ from choices under the first-best. Independent fishermen engage in inefficient search relative to the first-best.

More precisely, we see from panel (a) that the first-best policy depends on the sole owner's set of beliefs. For small levels of both μ_1/μ_2 and ν_2/ν_1 a single investment is warranted. The amount of investment under the first-best responds (weak) monotonically to changes in μ_1/μ_2 and ν_2/ν_1. That is, the number of fishermen that the sole owner sends to site 2 is nondecreasing in ν_2/ν_1 and nonincreasing in μ_1/μ_2. Panel (b) shows that in regions of the belief space where search by independent fishermen is optimal, two investments are made, and two search costs incurred. In this sense, the results are consistent with redundant search by independent fishermen over some region of the belief space. The results show that independent fishermen invest in information over a smaller region of initial beliefs than under the first-best.

Now consider the effects of congestion costs on actions. Non-cooperating fishermen play a mixed strategy over much of the belief space shown in the figure. Panel (d) shows the sum of probability weighted visits to site 2. From the perspective of the Row fisherman, with probability $\breve{a}_2$, his conjecture about the probability the Column player fishes site 2, an investment in information will come at an additional congestion cost. This reduces the incentive to resolve uncertainty at site 2; for any given investment cost, the threshold uncertainty level at which the investment is optimal must be higher. Through central coordination of fishing locations, congestion costs can be avoided. All else equal, the sole owner prefers to send fishermen to different sites and therefore, for a given investment cost the threshold level of uncertainty for which information investment is optimal is reduced. As in the case with no congestion costs the results show that, relative to the first-best, independent fishermen engage in redundant search over some regions of the belief space and underinvest in others. It should be noted that the congestion penalty alters the investment decision in such a way that the overall inefficiency is lessened. That is, within the region where fishermen play a mixed strategy, the probability that independent fishermen participate in redundant search is less than 1.

Result 3. (Inefficiency of independent search).

1. *Independent fishermen engage in inefficient search relative to the first-best. Independent fishermen only replicate the first-best search pattern when the value of information is at an extreme. That is, either when information has so little value that there is no investment under the first-best, or when it has a very large value so that the first-best invests with both fishermen.*

2. *A positive congestion* ($\kappa > 0$) *penalty lessens the inefficiency of independent fishermen by encouraging fishermen to play a coordination game.*

2.3.4 Risk Aversion

The previous analysis raises questions regarding how risk preference alters the demand for information, and whether independent fishermen respond to uncertainty differently than the first-best. Up to this point our analysis has been based on a linear utility specification. To highlight the role of information sharing, coordination, and risk preference on information demand, utility is parameterized to the negative exponential function.

$$U(s_j) = 1 - \exp(-\lambda s_j)$$

This specification of utility has several nice properties. First, this utility specification exhibits constant absolute risk aversion, where $\lambda > 0$ represents a fisherman's risk preference. Second, in the limit as $\lambda \to 0$ this utility specification is consistent with risk neutral preferences. Third, utility is bound to lie in the $(0, 1)$ interval, simplifying analysis. Finally, clean analytic expressions for expected utility exist. We begin with a discussion of the effect of changes in relative uncertainty on the decision to invest in information about site 2.

Under risk neutrality, the effect of changes in relative uncertainty is clear. While the cost of acquiring information is unaffected by a change, the expected gains from an investment are increasing in relative uncertainty due to the fisherman's ability to incorporate new information and reoptimize in the second fishing period. Under risk aversion, however, this partial effect is less clear. Though an increase in relative uncertainty still has no effect on the cost of an investment in information, given the disutility associated with selecting the less certain site, its effect on the expected gains from the information is ambiguous. Risk aversion imposes a "cost" on future choices through the disutility associated with unresolved uncertainty. For small levels of relative uncertainty, this creates a small downward pressure on the expected future gains of information. But as relative uncertainty grows, this downward pressure increases and the net gains can eventually become negative. As such, not only will it be optimal to search site 2 over a smaller region of belief space than under risk neutrality, but the relationship between relative mean payoffs and a threshold level of payoff uncertainty no longer yields the same monotonic tradeoff displayed in Figure A.2.

To illustrate, we return to a comparison of the equilibrium site choice strategies of independent fishermen with the first-best. Again, analytic solutions for these strategies are not available. To maintain comparability with the case of risk neutral preferences we calculate equilibrium site choices over the same range of parameter values. We fix the coefficient of absolute risk aversion to $\lambda = 0.1$.[13] Figure A.3 depicts our results, where the axes hold the same meaning as in Figure A.2.

There are a few features of the equilibrium search pattern, highlighted in Figure A.3, that deserve note. First, in juxtaposition with risk neutral preferences, both independent fishermen and the first-best explore site 2 over a smaller region of belief space. Also, due to the disutility of fishing an uncertain site, the first-best is now willing to invest, at most, a single vessel to fish site 2. Second, much like in the case of risk neutral preferences, independent fishermen still engage in inefficient search. As before, this inefficiency is lessened with the introduction of a congestion penalty, which encourages the fishermen to play a coordination game. Finally, the relationship between relative mean payoffs and a threshold level of payoff uncertainty is no longer monotonic. In fact, it exhibits a backward-bending relationship, capturing the increasing impact of disutility associated with experiencing uncertainty on the investment decision. This

[13] Our selection of λ does not qualitatively change our results.

backward-bending relationship also implies that there exists a maximum level of relative mean payoffs, above which no investment in information will occur.

Result 4. (Risk preference and the demand for information).

Consider the case where $\mu_1 > \mu_2$.

1. *Demand for information in site 2 is nonincreasing in the relative mean payoff.*

2. *Under risk neutrality, investment in information is nondecreasing in relative uncertainty while under risk aversion, it is nonmonotonic. That is, over lower levels of ν_2/ν_1 the demand for information is weakly increasing. As relative uncertainty increases, information demand peaks and then falls to zero.*

3. *Under information sharing and coordination, information demand is weakly increasing in the congestion penalty, while it is nonmonotonic under independent action as it promoted independent fishermen to play a mixed fishing strategy.*

4. *Investment in information is weakly decreasing in absolute risk aversion and occurs over a smaller region of belief space.*

2.3.5 Fishing cooperatives

In the first-best scenario considered above, fishermen were directed to sites by a fictitious sole owner, and no consideration was made to the internal distribution of payoffs. This construct ignores two crucial characteristics that distinguish cooperatives from owner-investor firms; the decision rights within the organization, and the residual claims on earnings. A comprehensive review of contracts that might be used to assign rights and residual claims within a fishing cooperative is beyond the scope of this paper.[14] We consider simple, but illustrative, examples of governance structures that might be operational in a real world fishing cooperative. Our approach follows the "new institutional economics" viewpoint which emphasizes the role of transactions costs, property rights and agency relationships for understanding the organizational structure (Fama (1980), Williamson (1975), and Alchian and Demsetz (1972)).

[14] See Fama and Jensen (1983) for additional discussion of contracting in producer organizations.

In the context of our fishing problem, the right to decide which site is fished by individual cooperative members will have important welfare and efficiency implications. In what follows we do not consider a formal assignment of decision rights. Instead, we will discuss, informally, the implications of simple decision rights that might exist within the fishing cooperative, such as majority or unanimous voting by members.

Let us return to the case of risk neutral preferences. Suppose cooperative members are compensated for their fishing efforts under a piece-meal remuneration contract. Such a contract design is common in fisheries, often used to promote truthtelling and avoid problems of effort shirking (McConnell and Price (2006)). Each fisherman receives a per-period payment, which we denote ω, plus a share, α, of the payoff realized at a fished site. Therefore, the expected single-period payoff from fishing site j is,[15]

$$\omega + \alpha\mu_j$$

Following some preliminary analysis we will consider a profit sharing arrangement where ω is determined by the aggregate payoffs earned by all cooperative members. This simple setup is general enough to capture egalitarian profit-sharing to full retention of the payoff at a site. We next show that a preference for exploitation over exploration will arise under a fisherman remuneration contract that allows fishermen to retain a disproportionate share of their realized fishing payoffs.

We assume that fishing cooperative members share common initial beliefs about true payoffs at competing sites. For consistency, suppose beliefs satisfy $\mu_1 > \mu_2$ and $\nu_1 < \nu_2$. The cooperative fishermen are capable of solving for the first-best site choice policy and agree that the first-best policy yields the highest expected payoffs for the cooperative. What is less clear is the mechanism that will determine where each cooperative member will fish, how payoffs are distributed, and whether the internal governance of the cooperative can implement the first-best policy.

No tension will arise if the first-best policy sends both fishermen to the same site. We therefore focus on the case where it is optimal, under the first-best, to send one fisherman to

[15] Under risk aversion, this contract decision has a natural risk sharing property.

each site. Suppose also that there is no congestion penalty; we fixate on the role of information. In expectation, $\kappa = 0$ implies that period 2 payoffs are identical for both fishermen; both will fish the same site, whichever has the highest expected payoff. We therefore focus on the expected payoff from the first period only. Lastly, suppose ω is used to redistribute the total earnings back to cooperating fishermen, and that the cooperative balances its budget in each period. Each member retains an α share of his own fishing payoff. The residual is re-distributed (equally) to cooperative members is $(1-\alpha)(s_1 + \check{s}_2)$.

For consistency, suppose beliefs satisfy $\mu_1 > \mu_2$ and $\nu_1 < \nu_2$. In expectation, a cost $\mu_1 - \mu_2$ must be incurred to obtain information about site 2's true payoff. Our assumption for profit redistribution implies an expected wage,

$$\omega = \frac{(1-\alpha)}{2}(\mu_1 + \mu_2).$$

Simple algebra finds that the first period payoff for the fisherman who exploits the relatively certain site 1 is,

$$\frac{1}{2}(\mu_1 + \mu_2) + \frac{\alpha}{2}(\mu_1 - \mu_2), \tag{2.14}$$

whereas, the payoff for the fisherman who explores the uncertain site 2 is,

$$\frac{1}{2}(\mu_1 + \mu_2) - \frac{\alpha}{2}(\mu_1 - \mu_2). \tag{2.15}$$

These expressions clearly illustrate that for $\alpha > 0$, a cooperative member will prefer exploitation over exploration.[16] Viewed another way, individual members will have an incentive to free-ride on the costly search efforts of the other cooperative fishermen.

A broader question is whether free-riding will impede the cooperative's ability to implement the first-best site choice policy? The issue here is the extent to which the internal distribution of information rents and costs distort the process that directs individual cooperative members to fishing sites.

To demonstrate, Figure A.4 shows the full expected payoffs under the first-best, and when remuneration follows a piece-meal contract with $\alpha > 0$. Specifically, the figure shows two-period

[16]This holds under risk aversion as well. Let $\pi_{12} = \frac{1}{2}(s_1 + \check{s}_2) + \frac{\alpha}{2}(s_1 - \check{s}_2)$ and $\pi_{12} = \frac{1}{2}(s_1 + \check{s}_2) - \frac{\alpha}{2}(s_1 - \check{s}_2)$. π_{12} and π_{21} are random variables. Using the convolution of s_1 and $\check{s}_2$ we can derive the mean and variance for these random payoffs and compute the corresponding expected utilities. Given our assumptions on beliefs, $E[\pi_{12}] > E[\pi_{21}]$ and $Var[\pi_{12}] < Var[\pi_{21}]$, which imply $EU(\pi_{12}) > EU(\pi_{21})$.

expected payoffs holding uncertainty fixed, with $\nu_1 < \nu_2$. Along the horizontal axis $\mu_1 - \mu_2$ increases, and therefore, represents increasing costs of gathering information at an uncertain fishing site, which as above is taken to be site 2.

Panel (a) in Figure A.4 reports the first-best average expected payoffs to the cooperative, denoted $\bar{V}_{11}$ and $\bar{V}_{12}$. It should be noted that rescaling the first-best payoffs does not change the site preference ordering. That is, it is equivalent to study the investment decision by examining total or average contribution to payoffs. To reduce clutter, $\bar{V}_{22}$ is not shown.

The first-best policy is determined by the maximum of $\bar{V}_{11}$ and $\bar{V}_{12}$. The point at which the cost and benefit of information are just equal is denoted as f in the figure. To the left of point f, it is optimal to send one fisherman to each site. Intuitively, when the cost of gathering information at site 2 is not too large it pays to gather an informative signal. When costs rise above f gathering information at uncertain site 2 does not pay and both fishermen fish at site 1.

Panel (b) depicts expected payoffs under a piece-meal contract with $\alpha > 0$. V_{12} is the payoff to the fisherman that exploits the relatively certain site 1 and V_{21} denotes the payoff to the fisherman that explores uncertain site 2. With $\alpha > 0$, $V_{12} > V_{21}$ and the exploring fisherman bears a larger share of the investment cost. Note also that the difference between the expected payoffs increases with α. Again, to reduce clutter, V_{22} is not shown.

Panel (b) illustrates the potential for internal disagreement over site choices. The exploiting fisherman benefits from an investment in information, and for mean payoff differences less than $\mu_1 - \mu_2 = g$ is likely to support a policy that sends some other fisherman to site 2. The fisherman assigned to explore the uncertain site will support a visit to site 2 only if the mean payoff difference is less than $\mu_1 - \mu_2 = e$. Further calculations reveal,

$$e = \frac{f}{1+\alpha} \quad \text{and} \quad g = \frac{f}{1-\alpha},$$

which demonstrate a clear trade-off between a performance-based remuneration and potential internal conflict over site choice. Whether this conflict leads to inefficient search will depend on the specific decision rights used to direct members to fishing sites and the size of the fishing cooperative. Suppose there are more than two fishermen in the cooperative. If, for example site choices are determined by a majority vote of cooperative fishermen, investments in information with costs in the range $[e, g]$ would be determined democratically; a potential over-investment in information between f and g. If, alternatively, all cooperative fishermen must agree on the site choice plan, i.e., only search costs below point e in Figure A.4 might be considered. In either case, site choices determined by voting rights can diverge from the first-best plan. Of course, this will result in lower total payoffs than under the first-best.

Can the piece-meal contract be modified to remedy the free-riding problem for cooperating fishermen? Suppose for example a fee is charged to fishermen who are the net beneficiaries of investments in information, and a subsidy is paid to fishermen who bear the burden of costly search. In order for fishermen to be indifferent between exploration and exploitation, the first-period remuneration must be equalized. From equations (2.14) and (2.15) we find that the fee collected from the exploiting fishermen and transferred to the exploring fishermen must fully offset the retained profits each fishing site. In other words, the fee and subsidy must counter the effect of the residual claim on realized fishing payoffs. A system of fees and subsidies that removes the free-riding problem and reproduces the first-best must correspond to the case of $\alpha = 0$, i.e., fishermen remuneration is independent of the realized fishing payoffs at the sites at which they fish.

Result 5. (Risk preference, congestion effects, and costly search).

1. *A member of the fishing cooperative prefers the fishing policy that maximizes his private expected utility. As such, he only incorporates his private gain and share of the cost of information into his investment decision.*

2. *A positive congestion penalty ($\kappa > 0$) reduces the cost of information, and thereby increases the region of belief space where the cooperative fishermen agree to invest in information.*

3. *Consider the case where $\mu_1 > \mu_2$. Given a governance structure, the fishing cooperative will replicate the first-best search pattern if $V_{11} > V_{12}$ or $V_{11} < V_{21}$. That if either the exploiting fisherman prefers both members fish site 1, or if the exploring fisherman prefers*

exploring site 2 rather than both fish site 1. In either case, there is no tension between the fishermen, despite any asymmetry in profit-sharing.

4. *Again, consider the case where $\mu_1 > \mu_2$. Given a governance structure for the cooperative, fishermen will disagree on the fishing policy if $V_{12} > V_{11} > V_{21}$. In this region of belief space the mechanism for selecting the cooperative's site choices will determine whether the fishing cooperative will replicate the first-best search pattern.*

5. *An incentive-based contract used to align incentives and prevent effort shirking under the profit-sharing agreement creates a free-riding problem on the costly investment in information.*

2.4 Conclusions

This paper analyzes search, learning and economic performance in a dynamic fishing game. We contrast equilibrium site choice strategies under a first-best benchmark, the case of independent, non-cooperative fishermen, and under a stylized fishing cooperative. In addition, we examine the effect of risk preference on equilibrium search patterns.

We find that relative to the sole owner benchmark, independent fishermen engage in redundant and inefficient search when the perceived cost of acquiring information is particularly low relative to the private value of the information. Redundant search by independent fishermen lessens as the penalty for congesting a site increases. Because independent fishermen do not share valuable information once it is acquired, they tend to under-invest in information relative to the first-best as the cost of information grows. Site choices of independent fishermen replicate the first-best only when the value of information relative to its cost of acquisition is at an extreme. The findings are driven by the fact that independent fishermen must bear the full cost of gathering information but do not benefit from the non-rival nature of information once acquired.

Our analysis shows that fishermen can benefit from sharing information about productive fishing sites and coordinating site choices to avoid congestion. Information regarding the productivity of a fishing location is a non-rival excludable good. This serves as both a benefit and problem for an information sharing group. While independent fishermen must privately bear the full cost of information, an information sharing group faces a potential free-rider problem. The cooperative must devise a method to distribute informational rents among its members, while maintaining incentives to undertake costly search for fish. The efficiency of a fishing co-operative depends critically on its internal governance structure. That is, how it addresses how rents are distributed, how cooperative site choices are determined, and how fishing locations are assigned among members. We demonstrate that under simple profit-sharing rules, common in fisheries, there are regions of the belief space where the private incentives of cooperative members conflict. The result can be inefficient search and reduced economic performance.

Fishing cooperatives represent a governance structure that can facilitate information sharing. An information sharing fishing cooperative can outperform independent fishermen, but only if contracts can be devised that provide members with correct incentives to undergo costly investments in information. Devising such contracts may be particularly challenging in fisheries. We show that eliminating free-riding problems may require the elimination of performance-based remuneration schemes. Egalitarian profit-sharing requires full disclosure, and likely costly monitoring of costs incurred by all cooperative members. Additional problems of truthtelling and shirking may arise without performance-based remuneration. It has been argued that agency problems in fisheries are pervasive, and their existence explains the almost universal use of revenue share contracts among skippers and crew members (McConnell and Price (2006)).

Maintaining sharp incentives to catch fish may take on added importance in commercial fisheries due to the common pool nature of fisheries resources. Regulations often force fishermen to play a constant sum game for limited total harvest. These management programs reward fishermen who can outperform rivals, and paying a skipper and crew a fixed wage in this setting may not be an option.

Risk preference is an important component of the problem of search with active learning. It affects both the characteristics of the demand for information and what can be learned through observation of the site choices of fishermen. Risk aversion introduces a nonmonotonicity in the demand for information with respect to relative uncertainty, which has consequences on what can be learned through observed behavior. To make accurate predictions regarding the spatial movement of fishermen and their response to fishery management, a measure of beliefs about payoffs at competing fishing locations and a measure of risk preference is required.

Extensions of the current model to formally model the membership decision under an endogenous contract design, and to examine the impact of strategic disclosure of information may provide additional insights. Deceit is common among commercial fishermen (Palmer (1990), Andersen (1980), and Andersen and Wadel (1972)). The practice of misinformation is likely to degrade the value of shared information and thereby weaken the value of a cooperative. Identifying conditions under which truthtelling is an equilibrium strategy in an information sharing game may provide additional insights and possibly suggest policies to improve performance in fisheries. An extension to consider cheap talk along the lines of Crawford and Sobel (1982) could provide further insights regarding information sharing in fisheries.

It is becoming increasingly evident that solutions to worldwide depletion of marine fish stocks will involve some form of strengthened property rights for resource users (Costello et al. (2008)). Allocating portions of sustainable harvest levels to a fishing cooperative, consistent with NOAA Catch Share Policy (National Oceanic and Atmospheric Administration (2010)), is a form of property right that promises to improve resource stewardship and enhance economic performance. Our results suggest that to capture the full benefits of information sharing and coordination, a fishing cooperative must address an internal governance problem that accompanies costly information acquisition. What is not clear is whether agency problems are more or less pronounced under alternative rights-based management approach, for example, a system of individual transferable fishing quotas. Further research on the relationship between regulations in the fishery, and information sharing and coordination would seem prudent.

Amendment 16 to the Northeast Multispecies Fishery Management Plan, developed by the New England Fishery Management Council, provides a recent case where a better understanding of the function and performance of fishing cooperatives is likely to pay dividends (New England Fishery Management Council (2010)). Under amendment 16, the National Marine Fisheries Service is responsible for determining the number of fishing cooperatives that conduct harvesting operations and the annual total harvests of a host of groundfish species. Approved sectors are allocated shares of the total groundfish harvest if they can demonstrate their ability to self-monitor cooperative members and ensure harvest allocations are not exceeded. This paper provides new insights concerning the internal governance of fishing cooperatives and raises questions for the overall efficacy of the sectors management approach. Particularly we show that free riding problems must be overcome in order to achieve efficient search within a fishing cooperative.

More generally, efficient fishery resource use involves selecting per-period harvest rates that maximize the social benefits from fishery exploitation. The impacts of alternative forms of regulation, e.g., input controls versus property rights-based approaches or sectors management have implications for information acquisition and therefore the costs of harvesting fish. Other fishery management problems related to bycatch of unwanted species, spatial closures, the design of marine protected areas have been linked to problems of information acquisition and information sharing in fisheries (Abbott and Wilen (2010), Haynie et al. (2009), Marcoul and Weninger (2008), Gilman et al. (2006), Costello and Deacon (2007), Curtis and McConnell (2004)). Our paper provides baseline results and a framework for further analysis in this important area of applied resource policy.

Bibliography

Abbott, J., Wilen, J. E., 2010. Voluntary Cooperation in the Commons? Evaluating the Sea State Program with Reduced Form and Structural Models. Land Economics 86 (1), 131–154.

Alchian, A., Demsetz, H., 1972. Production, Information Costs, and Economic Organization. American Economic Review 62 (5), 777–795.

Andersen, R., 1980. Secrecy: A Cross-Cultural Perspective. Human Science Press, Ch. Hunt and Conceal: Information Management in Newfoundland Deep-Sea Trawler Fishing, pp. 205–228.

Andersen, R., Wadel, C., 1972. North Atlantic Fishermen. University of Toronto Press, Ch. Hunt and Deceive: Information Management in Newfoundland Deep-Sea Trawler Fishing, pp. 120–140.

Buchanan, J., 1965. An Economic Theory of Clubs. Economica 32 (125), 1–14.

Carpenter, J., Seki, E., 2005. Do Social Preferences Increase Productivity? Field Experimental Evidence from Fishermen in Toyama Bay. IZA DP No. 1697, Institute for the Study of Labor (IZA).

Costello, C., Deacon, R., 2007. The Efficiency Gains from Fully Delineating Rights in an ITQ Fishery. Marine Resource Economics 22 (4), 347–361.

Costello, C., Gaines, S. D., Lynham, J., 2008. Can Catch Shares Prevent Fisheries Collapse? Science 321 (5896), 1678–1681.

Costello, C., Polasky, S., 2008. Optimal Harvesting of Stochastic Spatial Resources. Journal of Environmental Economics and Management 56 (1), 1–18.

Crawford, V., Sobel, J., 1982. Strategic Information Transmission. Econometrica 50 (6), 1431–1451.

Criddle, K., Macinko, S., 2000. A Requiem for the IFQ in U.S. Fisheries? Marine Policy 24, 461–469.

Curtis, R. E., McConnell, K. E., 2004. Incorporating Information and Expectations in Fishermen's Spatial Decisions. Marine Resource Economics 19, 131–143.

Deacon, R., Costello, C., Parker, D., 2008. A Model of Fishery Harvest with a Voluntary Cooperative, Department of Economics, UCSB Working Paper.

Fama, E. F., 1980. Agency Problems and the Theory of the Firm. Journal of Political Economy 88, 288–307.

Fama, E. F., Jensen, M. C., 1983. Separation of Ownership and Control. Journal of Law and Economics XXVI, 301–325.

Garland, W., Brown, P., May 1985. Fishery Cooperatives. Tech. rep., United States Department of Agriculture, ACS Research Report Number 44.

Gaspart, F., Seki, E., 2003. Cooperation, Status Seeking and Competitive Behaviour: Theory and Evidence. Journal of Economic Behavior and Organization 51 (1), 51–77.

Gatewood, J., 1984. Cooperation, Competition, and Synergy: Information-Sharing Groups Among Southeast Alaskan Salmon Seiners. American Ethnologist 11 (2), 350–370.

Gatewood, J., 1987. Information-Sharing Cliques and Information Networks. American Ethnologist 14 (4), 777–778.

Gelman, A., Carlin, J., Stern, H., Rubin, D., 2003. Bayesian Data Analysis, 2nd Edition. Chapman and Hall/CRC.

Gibbons, R., 1992. Game Theory for Applied Economists. Princeton University Press.

Gilman, E. L., Dalzell, P., Martin, S., 2006. Fleet Communication to Abate Fisheries Bycatch. Marine Policy 30, 360–366.

Haynie, A. C., Hicks, R. L., Schnier, K. E., 2009. Common Property, Information, and Cooperation: Commericial Fishing in the Bering Sea. Ecological Economics 69, 406–413.

Kitts, A., Edwards, S., 2003. Cooperatives in U.S. Fisheries: Realizing the Potential of the Fishermen's Collective Marketing Act. Marine Policy 27, 357–366.

Knapp, G., 2008. Case Studies in Fisheries Self-Governance. Ch. The Chignik Salmon Cooperative, pp. 335–348, FAO Fisheries Technical Paper No. 504.

Knapp, G., Hill, L., 2003. Effects of the Chignik Salmon Cooperative: What the Permit Holders Say. Tech. rep., University of Alaska Anchorage, UA Research Summary No. 1.

Lynham, J., 2006. Schools of Fishermen: A Theory of Information Sharing in Spatial Search, Department of Economics, UCSB Working Paper.

MacCall, A. D., 1990. Dynamic Geography of Marine Fish Populations. University of Washington Press.

Mangel, M., Clark, C., 1983. Uncertainty, Search, and Information in Fisheries. Journal of the International Council for the Exploration of the Seas 41 (1), 93–103.

Mangel, M., Clark, C., 1986. Search Theory in Natural Resource Modeling. Natural Resource Modeling 1, 1–54.

Marcoul, P., Weninger, Q., 2008. Search and Active Learning with Correlated Information: Empirical Evidence from Mid-Atlantic Clam Fishermen. Journal of Economic Dynamics and Control 32 (6), 1921–1948.

Matulich, S., Sever, M., Inaba, F., 2001. Fishery Cooperatives as an Alternative to ITQs: Implications of the American Fisheries Act. Marine Resource Economics 16, 1–16.

McConnell, K., Price, M., 2006. The Lay System in Commercial Fisheries: Origin and Implications. Journal of Environmental Economics and Management 51 (3), 295–307.

National Oceanic and Atmospheric Administration, 2010. National Oceanic and Atmospheric Administration (NOAA) Catch Share Policy. Available online: http://www.nmfs.noaa.gov/sfa/domes_fish/catchshare/index.htm.

New England Fishery Management Council, 2010. New England Fisheries Management Council (NEFMC): Final Amendment 16 to the Northeast Multispecies Fishery Management Plan. Available online: http://www.nefmc.org/nemulti/index.html.

North Pacific Fishery Management Council, 2011. (NPFMC) Discussion Paper on Cooperatives: Gulf of Alaska Chinook Salmon Bycatch. Available online at: http://www.fakr.noaa.gov/npfmc/current_issues/bycatch/211GOAChinookCoops.pdf.

Orbach, M. K., 1977. Hunters, Seamen, and Entrepreneurs : The Tuna Seinermen of San Diego. Berkeley: University of California Press.

Palmer, C. T., 1990. Telling the Truth (up to a point): Radio Communication among Maine Lobstermen. Human Organization 49 (2), 157–163.

Smith, M., 2000. Spatial Search and Fishing Location Choice: Methodological Challenges of Empirical Modeling. American Journal of Agricultural Economics 82 (5), 1198–1206.

Smith, V. L., 1968. Economics of Production from Natural Resources. American Economic Review 58, 409–431.

Stiles, R. G., 1972. North Atlantic Fishermen. St. Johns, Ch. Fishermen, Wives and Radios: Aspects of Communication in a Newfoundland Fishing Community, pp. 35–60.

U.S. Department of Agriculture, 2009. United States Department of Agriculture: Business and Cooperative Programs: Number of Cooperatives & Memberships by Major Business Activities 2009. Available online: http://www.rurdev.usda.gov/rbs/coops/data.htm.

Wilen, J. E., Richardson, E. J., 2008. Case Studies in Fisheries Self-Governance. Ch. Rent Generation in the Alaskan Pollock Conservation Cooperative, pp. 261–368, FAO Fisheries Technical Paper No. 504.

Williamson, O., 1975. Markets and Hierarchies: Analysis and Antitrust Implications. The Free Press.

Wilson, J., 1990. Fishing for Knowledge. Land Economics 66 (1), 12–29.

CHAPTER 3. FLEET RATIONALIZATION UNDER INDIVIDUAL TRANSFERABLE QUOTAS

A modified version of the paper to be submitted to a peer-reviewed journal

3.1 Introduction

Individual transferable quotas (ITQs) offer an effective tool to address over-capitalization and rent dissipation that otherwise plague traditional command and control management approaches. ITQs provide incentives to efficiently utilize available economies of scale and over time align fleet harvest capacity with target catch levels. This is a process commonly referred to as fleet rationalization. The long run efficiency gains that accompany rationalization are well-documented (National Research Council (1999), Squires et al. (1995), Arnason (1990b)), whereas the time required to reduce the fleet size, an understanding which vessels exit and which remain, the pricing of ITQs, and the efficiency implications during the transition are not well understood. Uncertainty surrounding fleet rationalization in a switch to a rights-based management system presents a significant hurdle to program designers and industry. Despite this, only a handful of studies attempt to analyze the period of transition from the pre- to post-ITQ long run equilibrium (exceptions include Vestergaard et al. (2005), Anderson (2000), and Weninger and Just (1997)).

The transition is an important consideration in the design and implementation of an ITQ management program. Matulich et al. (1996) note that during the transition to equilibrium "rents are generated and winners and losers are defined" (pp. 114). Many of the criticisms of ITQ management focus on this transition period.[1] For example, historic volatility of initial

[1]U.S. fishery management has also stated concerns regarding the effect of the transition on small fishery-dependent coastal communities (National Research Council (1999), Buck (1995)).

trading prices presents a potential obstacle to acceptance by industry (Anderson and Sutinen (2006, 2005), Larkin and Milon (2000)). There are several cases of this phenomenon across the world's ITQ managed fisheries. For example, Newell et al. (2002) find evidence of price volatility that diminishes over the first four years of the multispecies quota trading market in New Zealand. In addition to permit price volatility, experience with several U.S. fisheries finds that the transition from pre- to post-ITQ equilibrium conditions can take several years (Matulich et al. (1996)). Weninger and Just (1997), for example, document the slow transition to the cost-efficient fleet structure observed in the Mid-Atlantic Surf Clam and Ocean Quahog fishery. The authors find that four years after the introduction of ITQs, the clam fishery has not achieved the cost-efficient outcome. A slow fleet rationalization and price volatility suggest unrealized rents during the transition (Arnason (2002), National Research Council (1999)).

With an ongoing push for rights-based management in U.S. fisheries, e.g. the National Oceanic and Atmospheric Administration (2010) Catch Share Policy, understanding the economic fundamentals of the transition is increasingly relevant. This paper addresses three main questions: (1) what are the economic drivers that determine the rate of fleet downsizing, (2) what causes the observed quota price volatility during initial trading periods, and (3) what are the efficiency implications of a slow transition?

To address these questions, we develop a finite-horizon dynamic transition model and analyze the fleet size, configuration, and quota trading prices that emerge during the transition to the ITQ-regime fleet structure. The transition model allows us to characterize harvesting efficiency and rent redistribution along the transition path. In addition, we compare how the transition path responds to changes in key model parameters.

Our analysis shows how uncertainty over relative cost-efficiency translates into uncertainty over quota trading prices. As such, a component of the ITQ asset's market price is speculation over future trading prices. Given heterogeneous beliefs regarding relative cost efficiency, this results in volatile permit trading prices. As fishermen resolve uncertainty, speculation lessens and trading prices stabilize.

Heterogeneous uncertainty and learning also provide insight into observed patterns of vessel exit. Fishermen incorporate uncertainty, expectations over future trading prices, and costs of entry/reentry into their choke prices. These choke prices increase in the cost-efficiency of fishermen, such that at market trading prices the most cost-inefficient fishermen exit first. "Moderately" cost-inefficient fishermen may engage in delayed-exit strategies as there is value in waiting to resolve uncertainty. Our results show further how the optimal timing of individual exit, and exit patterns of the fleet are affected by the form and extent of initial uncertainty, as well as other model parameters. These findings also provide insight into ITQ program design. The near universal practice of allocating initial shares of quota based on historic pre-ITQ catch, promotes delayed-exit strategies and thereby prolongs the transition period during which the full efficiency benefits of ITQ management are unrealized.

The comparative dynamics results highlight the importance of key model parameters on the transition. For example, beliefs play a strong role on the transition, but not on the equilibrium. Increases in the average belief of fishermen about the mean cost-efficiency of the fishing population are shown to lower trading prices during the transition and slow the rate of vessel exit. For a "very" pessimistic fleet, one that overestimates the cost-efficiency, and therefore the economic performance of other fishermen, the transition is quick. Increases in other model parameters, such as the entry cost, time preference, and quota supply, provide similar results.

This paper is organized as follows. The next section provides a review of previous work on fleet downsizing and quota pricing in ITQ fisheries. Section 3.3 presents the fleet transition model used in the simulation analysis. Section 3.4 provides a discussion of the simulation results, including delayed-exit strategies, speculative bidding behavior, and the generation of economic rents during the transition. In addition, simulations are used to show how key model parameters, e.g. uncertainty, total quota levels, prices, etc., impact the price and fleet transition paths. Section 3.5 summarizes the main results and discusses further research.

3.2 Background and Previous Literature

It is well-recognized that ITQ management programs have the potential to correct many of the inefficiencies associated with both regulated and unregulated fisheries. By creating a tradable asset whose value is derived from cost-efficiency, the market, "as if by an invisible hand," transforms an over-capitalized, cost-inefficient fleet structure toward a fishery rent maximizing fleet structure (Christy (1973), Montgomery (1972)). ITQ programs are used to manage fisheries across the globe (see Costello et al. (2008), Grafton et al. (2006)). These programs offer case studies from which to learn about short run, transitionary, and long run effects. Canada, Iceland, and New Zealand use ITQ programs extensively to manage their fisheries. Since 1986, New Zealand's ITQ system has grown to cover 33 species and 150 markets for fishing quotas, while as of 1990, all of Iceland's fisheries are managed under ITQs; see Newell et al. (2002), Arnason (1990a). Currently, 10 U.S. fisheries, covering more than 28 fish species, are managed under ITQ management programs (National Oceanic and Atmospheric Administration (2011)).[2]

The economics literature has made many theoretical and empirical contributions to understanding the bioeconomic impacts of ITQ fisheries management. Theoretical work outlines the economic implications of ITQ-regime operating rules, fleet configuration and bioeconomic performance. Empirical studies on the effect of ITQ management fall within one of two types. Either an *ex post* comparison of ITQ-regime conditions with those observed under the previous management regime, or an *ex ante* study that predicts the equilibrium conditions expected to emerge under the ITQ-regime operating rules (Andersen et al. (2010), Lian et al. (2009), Newell et al. (2007), Weninger and Waters (2003), Newell et al. (2002), Dupont (2000), Grafton et al. (2000), Weninger (1998), and Lindner et al. (1992)). Brandt and McEvoy (2006), for example, attempt to predict the winners and losers from the introduction of an ITQ program to the Atlantic Herring fishery. Using a stochastic frontier analysis to estimate the distribution of cost-efficiency, they infer which ports, vessels by gear type, and vessels by relation to buyers will gain/lose after switching to the new management system.

Despite a generous amount of research on the long run efficiency and welfare effects of ITQ management, relatively few economists have studied the transition period. The few exceptions include Vestergaard et al. (2005), Anderson (2000), and Weninger and Just (2002, 1997)).[3]

[2]Western ITQ fisheries include Halibut and Sablefish, Bering Sea and Aleutian Islands Crab, Pacific Sabelfish, and Pacific Coast Groundfish. Eastern fisheries managed with ITQs include the Surf Clam and Ocean Quahog fishery, Golden Tilefish, Atlantic Sea Scallops, and Wreckfish. Grouper and Tilefish, and Red Snapper in the Gulf of Mexico are also managed by ITQs.

[3]The study of firm dynamics is not unique to fishery analysis, although the characteristics of fisheries provide a unique setting for its study. The I.O. literature provides a wealth of work on the subject. Deily et al. (2000), Dixit (1989), Fudenberg and Tirole (1986), and Ghemawat and Nalebuff (1985) provide a few example of the

Weninger and Just (1997) focus on the slow transition to the equilibrium after the introduction of an ITQ program. The authors argue that the introduction of the ITQ asset drives fleet downsizing. Uncertainty over stochastic quota prices encourage moderately cost-inefficient vessels to postpone exit and wait for a favorable trading price. Weninger and Just find that the opportunity cost of waiting is inversely related to cost-inefficiency, such that "very" cost-inefficient vessels exit first.

Weninger and Just (2002) extend the entry-exit model to include stochastic uncertainty over harvest costs. They demonstrate that this additional source of uncertainty lowers the trigger price for exit, thereby "dulling the incentive to abandon unproductive capital" (pp. 579).

Vestergaard et al. (2005) offer an alternative explanation for the slow transition to the equilibrium fleet structure. Rather than treating the ITQ as a valuable asset, they emphasize the role of sunk cost on the firm's entry-exit decision in a traditional deterministic investment model. The authors define the sunk cost as the difference between a unit of investment (e.g. for the fishing vessel) and its salvage value. Vestergaard et al. argue that the larger the sunk cost, the longer cost-inefficient vessels will remain active without reinvestment in capital. These "delayed-exit" fishers face increasing variable costs as their capital decays until cost-inefficiency prompts exit.

To gain insight into the functioning of ITQ markets, economists have constructed laboratory experiments intended to simulate permit trading behavior in an ITQ-managed fishery; see Moxnes (2007) and Anderson and Sutinen (2006, 2005). An important focus of this research is permit price volatility. Permit price volatility has been observed in fisheries adopting ITQs particularly in initial periods following program introduction, and is considered a shortcoming of rights-based management programs (Newell et al. (2002), Larkin and Milon (2000)). Anderson and Sutinen (2006) note that permit trading in the laboratory was carried out at prices that exceed the value of the permit implied under experimental design. The authors find that a restricting trading to the lease of fishing permits (prohibiting sale) during an initial period reduced price volatility in laboratory trading experiments.

Missing from the literature cited above is a structural analysis of the economic fundamentals

---

breadth of insights that have emerged as determinants of firm entry and exit dynamics.

which underlie the entry-exit decision and the pricing of the ITQ asset during the transition. The dynamic transition model presented next fills this gap. Fishermen have heterogeneous uncertainty regarding relative cost-efficiency and therefore quota trading prices. Endogenous trading prices reflect noisy information about the opportunity cost of holding quota. Fishermen resolve uncertainty over time through repeated interaction in the quota trading market. Our framework allows us to describe the price and fleet transition paths and derive insights on the effect of beliefs on fleet rationalization that are absent in earlier literature.

3.3 Model

Consider a fishery that adopts an ITQ management program at date $t = 1$. Fishery management allocates a fixed supply of quota, Q, equal to the pre-ITQ total allowable catch. Prior to $t = 1$, aggregate harvest is maintained through a system of entry restrictions and input-controls. For example, by imposing season length restrictions, per-trip catch limits, gear-use restrictions, or a combination of the three. The case we consider is one where the number of active vessels in the pre-ITQ fishery, denoted n_0, is excessive relative to the fleet size expected to remain once rationalization is complete.

Fishing vessels are owned by risk-neutral, profit maximizing fishermen and possess a variable returns to scale fishing technology.[4] We assume fishermen are heterogeneous with respect to their ability to catch fish. The heterogeneity of the population of fishermen is described by the distribution $\vartheta(d)$ on the support $d \in [0, \infty)$. Variation in ability manifests as differences in harvesting costs. Let $c(q; d)$ denote the cost to a fisherman with skill level d of harvesting q units of the fish stock. The harvest cost function is assumed increasing and strictly convex in harvest, and increasing in d. The parameter d can be interpreted as the "distance" to a cost frontier. High productivity levels correspond to low values of d. Active fishermen employ a single unit of capital each period. The per-period cost of employing this capital is $f > 0$. There exists a one-time sunk cost for entry/reentry into the fishery, $\psi \geq 0$; e.g costs associated

[4]A variable returns to scale technology exhibits many of the characteristics of real-world fishing technologies, and ensures a unique fleet size and configuration in the long run equilibrium. This choice is supported by empirical research (e.g., Weninger and Waters (2003), Dupont (2000), Weninger (1998)). The fleet structure that emerges in an ITQ-managed fishery exploits scale economies that are unrealized under input-controls.

with preparing a seaworthy vessel, finding a crew, etc. This technology is captured in following parameterization,

$$c(q;d) = \begin{cases} 0 & q = 0 \\ (1+d)q^2 + f & q > 0 \text{ and incumbent fisherman} \\ (1+d)q^2 + f + \psi & q > 0 \text{ and entrant fisherman} \end{cases}$$

Fishermen know their own productivity d but do not directly observe the productivity of other fishermen. The distribution describing the heterogeneity of fishermen is known up to a single parameter β, the mean skill level. Each fisherman holds uncertain beliefs about this parameter. Let $p(\beta; B)$ denote a fisherman's beliefs about the mean skill where B contains the mean and variance of these beliefs. Beliefs follow basic probability rules with $p(\beta; B) \geq 0$ for all $\beta \geq 0$ and $\int p(\beta; B) d\beta = 1$.

The ITQ trading price is determined by the market interaction of heterogenously skilled fishermen. As such, uncertainty over the distribution of skill translates into uncertainty over endogenous quota trading prices. The following analysis will focus on the impact of uncertainty over relative productivity on quota prices and entry/exit behavior. To simplify the analysis we assume a constant fish price, $p_t = p$, and fix $p < 2Q$. In addition, this will ensure that no single fisherman harvests the total allowable catch.

3.3.1 The Pre-ITQ Management Regime

In the pre-ITQ fishery, fishermen select the level of harvest that maximizes fishing profits given an exogenous market price for fish and regulations. The present value of earnings for a fisherman are given as,

$$\max_{q_t} \sum_{t=1}^{\infty} \delta^{t-1} \Big[pq_t - c(q_t; d+z) \Big] \tag{3.1}$$

where $\delta \in [0,1]$ is a discount factor common across fishermen, and $z \geq 0$ denotes added distance from frontier costs associated with harvesting under the pre-ITQ regulation.[5] The above expression illustrates two noteworthy features of the pre-ITQ fishery. First, payoffs do

[5]Limited entry and input-controls increase the per unit cost of harvesting. These management tools distort capital investment incentives in the race-for-fish setting, promoting substitution away from restricted inputs to a less efficient (unrestricted) mix (Munro and Scott (1985)). The term z incorporates the use of a potentially suboptimal mix of inputs into the fisherman's cost function as added distance from the frontier costs.

not directly depend on the skill levels and capital costs of other fishermen.[6] To control for any indirect connection, we assume that the pre-ITQ fishery is in a long run equilibrium prior to the switch to ITQ management. Therefore, uncertainty over the distribution of skill does not affect the harvest decision, and is not relevant under the pre-ITQ management regime.

Second, under our model assumptions, harvest choices can be made independently over time; the problem in (3.1) is not dynamic. Define $q_t^* = \arg\max\{pq_t - c(q_t; d + z)\}$. In period t, a fisherman harvests q_t^* if variable profits exceed f. Otherwise, it is optimal to harvest zero and avoid the fixed cost f. Failure to cover capital costs suggests the fisherman allocates the capital embodied in the vessel and skill to its next highest-valued use, in other words, exit the fishery.

3.3.2 ITQ Management Regime

Under an ITQ management program, the per-period harvest of each fisherman must be matched by quota. The defining feature of the management approach is that quota is tradable. Assuming a well-functioning trading market exits, ITQs introduce a user cost that is absent under the pre-ITQ regime (Weninger and Just (1997)). In the following, the quota trading price is initially treated as a parameter of the model, used to characterize quota demand. The model and this section are closed with a description of the solution of the equilibrium price and formation of beliefs about future trading prices.

To simplify analysis of the ITQ fishery, each fishing period $t = 1, \ldots, T$ is separated into two activities, quota trading followed by fishing. With a slight abuse of notation, q will denote both the harvest level and quota holding of fishermen in the ITQ fishery. Given each fisherman's endowment of quota q_{t-1} and beliefs about the heterogeneity of fishermen B_{t-1}, fishermen interact in the quota market. Fishermen bid and offer for quota until a market clearing price r_t emerges, at which trades occur. The interaction of fishermen provides information about the true distribution of skill in the form of a public signal s_t. Fishermen use this information to update beliefs B_t following a prescribed learning rule common to fishermen; see Appendix B.1

[6] Regulations under an input-control management program respond to changes in aggregate fishing pressure, which is determined in part by the productivity of the entire fishing fleet. This creates an *indirect* dependence between skill levels and the long run profits of a fisherman.

for a description of the learning rule and public signal. Each fishing period ends with fishermen harvesting their updated quota holding, q_t, which is assumed private information. The timing of events is illustrated in Figure B.1.

An ITQ is a claim on a stream of profits. Let $\pi(q;d) = pq - c(q;d)$ denote the per-period profit of a fisherman with skill d harvesting q units of the fish stock. Each period, fishermen adjust their quota holding to maximize the expected discounted stream of profits.[7] Let x_t denote the quantity of quota bought/sold at the quota trading price r_t in period t. A fisherman's adjustment of quota is constrained by his current quota holding q_{t-1} and the fixed supply Q; $q_t = q_{t-1} + x_t$ where $x_t \in [-q_{t-1}, Q - q_{t-1}]$.

In the ITQ fishery, fishermen select the level of harvest/quota that maximizes fishing profits given exogenous market price for fish, the quota trading price, and expectations. The present value of earnings for a fisherman are given as,

$$\max_{q_t} E\Big\{\sum_{t=1}^{T} \delta^{t-1}\Big[p_t q_t - c(q_t;d) - r_t(q_t - q_{t-1})\Big]\Big|B_0\Big\}. \tag{3.2}$$

The above expression illustrates three noteworthy features of the ITQ fishery. First, under the ITQ management regime, fishermen are free to choose the optimal (cost-efficient) mix of inputs for harvesting the fish stock; $z = 0$. Asche et al. (2009) attribute this readjustment of inputs as one of the sources of rent generation from implementing an ITQ management program.

A second feature of ITQ fishing behavior is the sharpened dependence of payoffs on the skill level of other fishermen. The quota trading price is a market clearing price, which depends on the heterogeneity of fishermen. Expectations play a strong role in quota adjustment and the trading price that emerges each period.

Finally, unlike the pre-ITQ fishery, harvest choices are dynamic. The quota price changes over time in response to changing expectations. This creates the potential for capital gains and losses by fishermen, which is a key force in forward-looking behavior. The past actions affect the cost of quota adjustment in the current period, while expectations over future prices determine the level of this adjustment. As noted by Weninger and Just (1997), in this setting

[7] The finite horizon can be extended to an infinite horizon problem by allowing $T \to \infty$. Additional insight is not gained in the infinite horizon specification as the focus of this paper is the behavior of fishermen during the transition.

the entry/exit criterion does not follow standard Marshallian principles.

The Bellman equation for the fisherman's optimization problem has the following form,

$$V_{t-1}(q_{t-1}, B_{t-1}) = \max_{0\leq q_t\leq Q} \left\{\pi(q_t; d) - r_t(q_t - q_{t-1}) + \delta E\left[V_t(q_t, B_t)\middle|B_{t-1}\right]\right\} \tag{3.3}$$

for any period $t \leq T$ with $V_T = 0$. Given an initial allocation of quota and beliefs the solution to equation (3.3) determines a fisherman's optimal quota holding. The above dynamic programming problem contains two possible corner solutions, $q_t = 0$ (exit the fishery) and $q_t = Q$ (monopolize the fishery), along with the standard interior solution, $q_t \in (0, Q)$. We can ignore the corner solution $q_t = Q$ given the assumption on the price of fish; $p < 2Q$.

The Karusch-Kuhn-Tucker (KKT) conditions provide insight into the optimal behavior of a fisherman. From the problem in (3.3), these conditions are

$$\begin{aligned} &\pi_q(q_t; d) + \delta E\left[r_{t+1}\middle|B_{t-1}\right] \leq r_t, \\ &q_t \geq 0, \\ &q_t\Big(\pi_q(q_t; d) + \delta E\left[r_{t+1}\middle|B_{t-1}\right] - r_t\Big) = 0. \end{aligned} \tag{3.4}$$

where $\pi_q(\cdot)$ denotes the marginal profit.[8] The complementary slackness condition contains the three possible solutions to the optimal quota holding problem. First, the first-order condition could hold with equality implying that the optimal quota holding in period t equates the fisherman's marginal profit plus the expected discounted market value for a unit of quota in the next period to the current quota trading price. Second, the first-order condition could be a strict inequality implying that the optimal quota holding is zero, as the marginal gain in profits for any positive holding is less than its marginal cost. Third, the first order condition could hold with equality at a quota holding of zero. This last case can be ignored because of the U-shaped average costs.

3.3.3 Incumbent and Entrant Fishermen

To account for the corner solution in the optimization problem, "trigger prices" are needed to characterize the demand for quota. The inclusion of an entry fee affects the trigger prices of incumbent and entrant fishermen differently. In the following these trigger prices are described.

[8]Using backward induction, $\partial E\left[V_t(q_t, B_t)|B_{t-1}\right]/\partial q_t$ equals the discounted expected quota trading price in the next trading period.

In period t, an incumbent fisherman will remain active if and only if the value of remaining active (which includes the option to delay exit until a future period) is greater than or equal to the value of exit (with the option to reenter the fishery). Similarly a fisherman will enter the fishery if and only if the value of entry (including the cost of entry with the option to exit in the future) is greater than or equal to the value of remaining outside the fishery (with the option to delay entry). Then each fisherman's "choke price" equals the largest quota trading price such that the fisherman just prefers being active in the fishery. Let $\bar{r}_t$ denote the choke price in period t. Manipulation of equation (3.3) provides,[9]

$$\bar{r}_t(q_{t-1}) = \begin{cases} p - 2\sqrt{(f - \delta\Lambda_t)(1+d)} + \delta E[r_{t+1}|B_{t-1}] & \text{if } q_{t-1} > 0 \quad \text{(incumbent)} \\ p - 2\sqrt{(f + \psi - \delta\Lambda_t)(1+d)} + \delta E[r_{t+1}|B_{t-1}] & \text{if } q_{t-1} = 0 \quad \text{(entrant)} \end{cases} \tag{3.5}$$

where,

$$\Lambda_t = E\left[V_t(q_t, B_t)|B_{t-1}\right] - E\left[V_t(0, B_t)|B_{t-1}\right] - E[r_{t+1}|B_{t-1}]q_t.$$

Λ_t denotes the net expected value of remaining active in period t less the expected capital gains from holding quota. As such, this term captures the probability weighted payoffs of responding optimally to market trading prices in future periods, including the cost of reentry.

The form of the choke price in equation (3.5) generates some important insights. First, so long as there is an entry cost, the choke price for an incumbent fisherman with skill d exceeds the choke price for an entrant fisherman with the same skill level, creating a hysteresis gap. For the incumbent fisherman, the entry cost introduces a penalty for exiting the fishery "too soon," while it represents a barrier to entry for an entrant. As such, ψ raises the choke price of the incumbent but lowers the choke price of the entrant thereby widening the hysteresis gap. Finally, if there is no cost to entry, the choke prices of the incumbent and entrant fisherman coincide and the hysteresis gap vanishes.[10]

Second, expectations about future quota trading prices enter directly into choke prices. While the expected trading price for the next fishing period is transparent in equation (3.5), the

[9] These choke prices are consistent with Dixit (1989). The author finds that trigger prices for entry/exit in a dynamic and uncertain model will exceed the deterministic Marshallian trigger prices. As such, firms will enter and exit "later" than a deterministic Marshallian model would predict.

[10] Λ_t has an analytic solution and is solved through backward induction. If $\psi = 0$, then $\Lambda_t = 0$ and $\bar{r}_t = p - 2\sqrt{f(1+d)} + \delta E[r_{t+1}|B_{t-1}]$.

expected, discounted stream of trading prices over the remaining life of the asset are embedded in Λ_t. As such, choke prices are forward-looking over the remaining life of the asset. If the quota trading price is expected to rise in future periods, a fisherman is willing to accept a higher trading price in the current period to capture the expected capital gains.

Finally, choke prices are decreasing in skill. Less efficient fishermen require a lower price to ensure that participation is profitable.

Exploiting the KKT conditions and the choke prices in (3.4) and (3.5), defines a contingent strategy profile for the optimal quota demand in period t as a function of the quota trading price, endowment of quota, beliefs, and the fisherman's cost efficiency.

$$q_t = \begin{cases} 0 & \text{if } r_t > \bar{r}_t(q_{t-1}) \\ \big(p - r_t + \delta E\big[r_{t+1}\big|B_{t-1}\big]\big)/\big(2(1+d)\big) & \text{if } r_t \leq \bar{r}_t(q_{t-1}) \end{cases} \tag{3.6}$$

As expected, quota demand is decreasing in the quota trading price and increasing in the price of fish. Also, higher skilled fishermen harvest a larger quantity of the fish stock, consistent with the efficiency gains associated with tradable property rights.

More interesting is the speculative effect on quota demand. The discounted, expected quota trading price in the next period enters directly into demand. If a fisherman expects prices to rise in the future, he is willing to purchase more quota today in expectation of enjoying capital gains in the future. Similarly, if he expects prices to fall, he will delay purchase. Given endogenous trading prices, this is an avenue for self-fulfilling, or bubble behavior, that may lead to inflation of the market price.

Finally, this strategy profile has implications on the profits of incumbent fishermen. At a sufficiently high quota trading price an incumbent will sell his quota holding and exit the fishery. This implies there exists a minimum profit, which occurs at the incumbent fisherman's choke price.[11]

3.3.4 Market Clearing and Price Expectations

Up to this point the quota trading price has been treated as a parameter of the model. However, it is the interaction of heterogeneous fishermen that determines the market clearing

[11] Weninger and Just (1997) find similar results.

price. This price equates, as close as possible, the fixed supply of quota to aggregate demand.[12] Given the contingent strategy profiles of incumbent and entrant fishermen, this requires that in period t,

$$\sum\nolimits_i q_t\left(p, r_t, d_i, B_{i,t-1}\right) \leq Q \tag{3.7}$$

where i indexes fishermen in the population.

Quota demand, depicted in equation (3.7), depends on beliefs about the stream of future prices. Fishermen use the market clearing mechanism to form these beliefs. While the mechanism is common knowledge, fishermen do not know the private information embedded in each fisherman's quota demand; e.g. every fisherman's skill, beliefs, and quota holding. Instead, fishermen use their beliefs about the distribution of skill to infer a set of beliefs about the remaining stream of market clearing prices.[13] Given a fisherman's uncertainty over the distribution of skill, the remaining $T-t$ expected market clearing conditions are,

$$\int \sum\nolimits_i q_\tau\Big(p, r_\tau, x, B_{t-1}\Big)\vartheta(x|\beta)dx = Q \quad \text{for } \tau = t+1,\ \ldots,\ T \tag{3.8}$$

Beliefs are derived recursively. Starting with $\tau = T$, for a given β a fisherman can solve for the unique price that satisfies equation (3.8). This price is assigned a probability according to his beliefs, $p(\beta|B_{t-1})$. Then iterating across the possible values of β, this maps out a distribution of beliefs for the market price in period T. These beliefs are used to find beliefs for the trading price in period $T-1$, and so forth until $\tau = t+1$.

3.4 Results

Analytic solutions to the dynamic programming problem in equation (3.3) are not available. This section describes the results for a series of simulation exercises. A baseline simulation is used to highlight the implications of the model regarding delayed-exit strategies, speculation, and the generation of rents. Comparative dynamics on the baseline model finish this section. These simulations illustrate the impact of key model parameters on the transition.

[12]Aggregate demand is downward sloping, jumping discontinuously at choke prices. The market price minimizes excess demand such that aggregate demand does not exceed the supply of quota.

[13]For tractability, fishermen believe that they share the same beliefs and that the "average" fisherman has a positive endowment of quota. More realistic assumptions regarding *beliefs about beliefs* muddy the analysis without creating new insight.

In the baseline simulation, the heterogeneity of skill in the fishing population is described by an exponential distribution with inverse scale equal to 2 (mean productivity of $d = 0.5$). To preserve this distribution in the following analysis, skill levels are fixed across simulations. The interval $(0, 1)$ is divided into 50 evenly spaced points, and evaluated at these points using the inverse exponential distribution. This generates 50 fisherman productivity levels consistent with the distribution; productivity levels range from $d = 0.009$ to $d = 1.966$.

Fishermen are modeled as having heterogeneous beliefs about the mean skill. Beliefs follow a gamma distribution with mean values that range from 0.75 to 1.25 (variance between 0 and 1). As such, the simulations focus on an "optimistic" fleet of fishermen; e.g. fishermen believe that their productivity relative to the average is higher than it actually is.

Fishermen have a constant discount rate of 0.9, face an ex-vessel price of fish of \$1, a per-period fixed cost of \$0.1, and a one-time sunk of \$0.1 for entry into the fishery. The total allowable catch is set equal to $Q = 10$. Fishermen are initially allocated a portion of the TAC proportionate to historic catch in the pre-ITQ fishery. To gain insight into the dynamic behavior of fishermen while maintaining tractability of the model, T is set equal to 3. Finally, to control for the random assignment of beliefs, each simulation is run 200 times. The results presented below, e.g. quota trading prices, fleet size, etc., denote average values across runs of the simulation.

3.4.1 The Pre-ITQ Fishery

The parameter values from the baseline simulation result in 44.0 active fishermen, each earning an operating profit close to zero, and harvesting on the increasing returns to scale segment of their average cost curve. Active fishermen range in skill from $d = 0.009$ to $d = 0.993$, effectively 99.3% above frontier cost.

Table B.1 provides a description of the fleet configuration, harvest, and rents in the pre-ITQ fishery. Heterogeneity in the productivity of fishermen allow active fishermen to earn positive profits as a return to skill. Scale inefficiencies, participation of inefficient fishermen, and input regulations dissipate resource rents. The median active fisherman operates at 29.1% above frontier cost, harvests 0.2303 units of the TAC, and earns \$0.0427 in profits.

3.4.2 The ITQ Fishery

Rationalization of the fleet is not immediate. Tables B.2 and B.3 describe the transition to the new long run equilibrium. In the initial period of the ITQ trading program, quota trades at an average price of \$0.6836 per unit. At this price, 10.0% of the fishermen active in the pre-ITQ fishery choose to sell their endowment of quota and exit, leaving an active fleet of 39.6 fishermen. The fishermen that exit the fishery are the least skilled, improving the overall efficiency of the fleet. The median fisherman in the fleet is 12.8% closer to frontier costs. The market redistributes quota from less efficient fishermen to the more skilled fleet, raising total rents in the fishery by 60.6%.

In the second period, the quota trading price drops to an average value of \$0.4492 per unit. The ownership price for quota changes for two reasons. First, there are fewer trading periods remaining in the finite horizon game and therefore fewer periods to either use the quota to harvest fish or resell it for a capital gain. This naturally lowers its value. Second, some uncertainty has been resolved. This alters each fisherman's demand for quota as well as their choke prices, thereby affecting the market clearing quota trading price. At the new quota trading price, 15.6% of the remaining active fishermen sell their quota holding and exit the fishery, leaving an average fleet size of 33.4 fishermen. Again, the fishermen that exit are the least efficient among the active fishermen, leading to a redistribution of quota to the more skilled and raising total fishery rents by another 3.7%. The cost efficiency of the median fisherman in the fleet improves by 18.8%, relative to period 1.

In the final trading period, the trading price converges to an average long run equilibrium price of \$0.2312 per unit. At the long run price, another 7.5% of the fleet exits the fishery and the corresponding redistribution of quota raises total rents by another 2.6%.

The baseline simulation reveals that as predicted in the theoretical literature on ITQ fisheries, a more efficient fleet prevails under the ITQ management approach. In the long run equilibrium 31.0 fishermen harvest the total allowable catch, a 29.7% reduction in the size of the fleet. The rationalized fleet is more cost-efficient at harvesting the fish stock, with skill levels ranging from $d = 0.009$ to $d = 0.468$. This represents a 52.9% improvement in the cost efficiency of the least efficient fisherman in the fleet. Total fishery rents have also increased by 70.9% through removal of excess capital and effort, and a redistribution of harvest responsibilities to the most cost-efficient members of the population. The median active fisherman in the long run equilibrium operates at 18.8% above frontier cost, harvests 0.3237 units of the TAC, and earns $0.0992 in profits.

3.4.3 Delayed-exit

There is value in waiting. Uncertainty by fishermen regarding the opportunity cost of holding quota, along with costly reentry, create value in delaying exit to learn more information. Such a delayed-exit strategy may be profitable on moderately cost-inefficient vessels and therefore slow fleet downsizing. The results in the baseline simulation provide evidence of this behavior.

Weninger and Just (1997) define the opportunity cost of holding quota on a cost-inefficient vessel as the value that same unit of quota could provide on a cost-efficient vessel. The discrepancy in values across fishermen generate incentives to trade. But given uncertainty over the productivity, and therefore profitability of the fishing fleet, fishermen are also uncertain about the opportunity cost of holding the ITQ asset. Each period, some uncertainty is resolved through interaction in the quota trading market. Because reentry is costly, exiting “too soon” comes at a penalty. For some uncertain cost-inefficient fishermen, delaying exit and learning is more profitable than exiting with the possibility of reentry at a future date.

Fishermen incorporate uncertainty, expectations over future trading prices, and the cost of entry/reentry into their trigger prices, so that endogenous quota trading prices signal entry/exit to fishermen.[14] As noted earlier, trigger prices are strictly decreasing in skill, which suggests a pattern of exit. The most cost-inefficient fishermen will have the lowest trigger prices and exit first, while moderately cost-inefficient vessels will postpone exit.

To illustrate, Figure B.2 depicts the quota holding of heterogeneous fishermen for a single run of the baseline simulation. The horizonal axis captures the varying cost efficiency of

[14]Price signals are an important component of fleet rationalization (Newell et al. (2007)) and management of ITQ fisheries (Batstone and Sharp (2003),Arnason (1990b)).

fishermen, and the vertical axis their corresponding holding of quota. The dark dashed line in the figure represents the initial allocation of quota (q_0). As this is based on historic catch, the most productive fishermen are assigned the largest endowment. However, this also implies that cost-inefficient fishermen operating in the pre-ITQ fishery are assigned quota. The light solid line depicts the redistribution of quota in the long run equilibrium (q_3^*). In comparison to the initial endowment, quota is redistributed from the least efficient fishermen to the most. The transition from q_0 to q_3^* is not instantaneous. Instead, the figure illustrates a slow transition, with the least efficient fishermen exiting first. As moderately cost-inefficient fishermen learn about the opportunity cost of holding quota, they exit the fishery and their quota holding gravitates to the hands of the more efficient fisherman. This suggests a delayed-exit strategy.

This is an important result for policy makers. It is a near universal practice of U.S. fishery management to allocate the initial endowment of quota based on historic catch in the pre-ITQ fishery (National Research Council (1999)). However, this means that cost-inefficient vessels are endowed with a positive quota holding and may find delayed-exit strategies profitable. This implies, that contrary to conventional wisdom, the initial allocation of quota has efficiency implications, as delays in the transition to equilibrium represent foregone rents in the fishery.

3.4.4 Speculation

A component of the ITQ asset's price, during the transition, is speculation. This result is best understood by separating the fisherman's *fundamental* valuation of quota from the speculative value. Following Morris (1996) and Harrison and Kreps (1978) suppose for a moment that the fisherman could not resell the quota asset. In this case, its value to the fisherman would be determined solely by the present value of the stream of operating profits that could be earned from fishing the quota. The speculative value arises from the option to sell the asset to some fisherman that values it more. The speculative premium is therefore non-negative as the option to sell will be exercised only if the fishermen benefit from doing so.

Uncertainty about the distribution of productivity and therefore the quota asset's value to other potential buyers drives the speculative premium. Therefore beliefs about the distribution of productivity have an important effect on the quota price and fleet rationalization. Over time, as fishermen learn, beliefs converge to the true distribution, the speculative premium diminishes and the quota trading price converges toward the fundamental valuation held by the active fleet.[15]

This is consistent with the simulation results presented above. Speculative behavior is suggested by the market pricing of the ITQ asset. Consider the average period 1 quota trading price from the simulation, \$0.6836. Converting this ownership price into a per period equivalent price we find \$0.6836/2.71 = \$0.2522, which is 9.1% above the average long run price. In period 2, fishermen have resolved some uncertainty about the heterogeneity of fishermen; e.g. learned information about the opportunity cost of holding quota. As such, the speculative premium in the pricing of the ITQ asset should be smaller. Again, converting the ownership price into a per period equivalent, we find \$0.4492/1.9 = \$0.2364, which is 2.2% above the average long run price. One should exercise caution in crediting all of the markup as the speculative premium. Identification of this premium is confounded by the presence of a cost to entry. Fishermen incorporate the one-time sunk entry/reentery cost into their choke prices, requiring a higher quota trade price to induce exit. As such, the market may clear at a price above the long run equilibrium to account for this in the presence of uncertainty.

3.4.5 Economics Rents

Consistent with the theoretical literature, the ITQ program generates rents in the fishery. Under the pre-ITQ program there are two potential sources of rents; rents to the resource and rents to skill. Input-controls and excess effort of fishermen in pursuit of perceived economic opportunities dissipate resource rents. All remaining profits to the fishery are attributed to the heterogeneity of skill of the fishing fleet. That is, cost-efficient fishermen receive a premium that is inversely related to their distance from frontier costs.[16] This premium converges to zero as we move to the least efficient fisherman in the fleet. In the baseline simulation, presented above, rents under the pre-ITQ management regime total \$1.7890.

[15] In the long run equilibrium, the quota trading price equates the marginal profits of the fleet. Therefore, the fundamental valuations of active fishermen will also be equated.

[16] See Colgan and Pascoe (1999) and Anderson (1989) for a deeper discussion on the decomposition of rents in a managed fishery.

The ITQ fishery generates rents through the removal of redundant capital and effort, and the redistribution of the total allowable catch to the most cost-efficient vessels. In the long run equilibrium, the quota trading price represents the per unit return to the resource and equates the marginal profits of the fleet. Resource rents from the fully rationalized fleet equal the market value of the supply of quota, \$2.312. Remaining economic profits are attributed to returns to skill, totaling \$0.7459. Combining these two sources generates total fishery rents equal to \$3.0579, 70.1% above those generated in the pre-ITQ fishery.

Decomposition of rents during the transition is confounded by the effects of uncertainty, speculation, etc., all of which affect the endogenous quota trading price (Lindner et al. (1992)). Manipulation of the complementary slackness condition allows a separation of the resource rent in period t from speculation over future capital gains.[17]

$$r_t Q = \sum_i \pi_q(q_{it}) q_{it} + \delta \sum_i E\left[r_{t+1} \middle| B_{i,t-1}\right] q_{it}$$

The left-hand side of the equation represents the market value of the quota supply in period t. At the market trading price, this must equal the return to the fleet from restricting aggregate harvest (to the TAC) plus the discounted, expected market value of the quota supply in the next trading period (the discounted sum of expected capital gains from asset resale in period $t+1$). The first term on the right-hand side is proportional to the resource rents in period t.[18] The second term on the right-hand side captures the effect of uncertainty and speculation on the future value of quota. Consistent with the pre-ITQ calculation, all remaining economic profits in excess of the market value of quota (reweighted to a single period value) are attributed to the skill premium of heterogeneous fishermen. Table B.4 describes the composition of rents during the transition.

In the initial trading period, 58.7% of economic profits in the fishery are attributed to resource and skill rents. As uncertainty is resolved, cost-inefficient fishermen sell their quota holding and exit the fishery. The redistribution of harvest responsibilities to more cost-efficient vessels allow these fishermen to enjoy scale economies. This improves both the resource rent

[17] The complementary slackness condition can be rewritten as, $r_t q_{it} = \pi_q(q_{it}) q_{it} + \delta E[r_{t+1}|B_{it-1}] q_{it}$. Summing over all fishermen provides the above equation.

[18] The resource rent in period t is calculated using the scaling constant $\omega_t = 1 \Big/ \sum_{k=t}^{3} \delta^{k-t}$.

and the returns to skill during the transition, while the effect speculation decreases. In the long run equilibrium, these are the only sources of economic rents.

The majority of rents in the long run equilibrium are generated in the initial trading period (85.5%), through the removal of excess capital, relaxation of input-controls, and the redistribution of harvest responsibilities. In each subsequent period additional rents are generated as rationalization continues. Over the transition, 14.5% of long run rents were lost due to delayed-exit and uncertainty.[19]

3.4.6 Comparative Dynamics

Selection of the model parameters is an important consideration in the simulation analysis. This choice can affect both the transition to and determination of the long run equilibrium. In the following, we describe a series of comparative dynamics exercises to examine the impact of model parameters on the transition path, and test the sensitivity of results from the baseline simulation. Table B.5 contains a summary description of the findings for key model parameters.[20]

Consider first the effect of the average mean belief in the fishing population.[21] A fisherman's mean belief about relative cost-efficiency affects his expectations about future trading prices, quota demand, and choke prices. This parameter is important in the determination of market trading prices and the rate of vessel exit. To explore its effect, average mean beliefs are varied from 0.25 to 3.00 across simulations. This covers a range of beliefs about relative economic performance. Table B.6 and Figure B.3 contain these results. The results suggest that a fishery comprised of fishermen that overestimate the cost-efficiency of other fishermen, illustrated by $\bar{\mu} = 0.25$ in Figure B.3, will rationalize quicker and experience higher trading prices during the transition. These fishermen overestimate the opportunity cost of holding capital and wish to sell their quota holding and exit the fishery. As fishermen, on average, expect higher trading

[19]The above analysis considers the comparison of pre-ITQ rents with the long run equilibrium. There are different ways to calculate the loss of economic rents during the transition. For example, the loss in the discounted stream of rents is only 3.1%. If we only consider the sustainable portions of rent (in equilibrium), then the loss in the discounted stream of rents is much higher, 27.1%.

[20]All results are relative to the baseline simulation.

[21]As beliefs are heterogeneous, the average mean belief refers the average of the heterogeneous beliefs of fishermen about the mean skill level. The range of heterogeneous mean beliefs is held constant across simulations.

prices this pushes up the market trading price. The reverse case holds similarly. If fishermen, on average, underestimate the cost-efficiency of other fishermen, as illustrated by $\bar{\mu} = 3.00$ in the figure, the rate of vessel exit slows and trading prices fall.

The entry cost represents a barrier to entry. Changes in this cost alter choke prices and therefore affect both the rate of vessel exit and quota trading prices during the transition. To explore this effect, the entry cost is varied from 0.00 to 0.25 in a series of simulations. Table B.7 and Figure B.4 contain the average simulation results for these values. The entry cost enters directly into the choke prices that define the entry-exit decisions of fishermen (see equation (3.5)). If there is no cost to entry ($\psi = 0$), this lowers choke prices and eliminates the hysteresis gap. In this case we observe many vessels exiting immediately. Depending of the beliefs of fishermen too many/few exit initially and further adjustments are necessary. As the cost of entry increases, so do choke prices, encouraging moderately cost-inefficient fishermen to engage in a delayed-exit strategy.

The per period fixed cost of operating capital is important in determining the minimum efficient scale of harvest, the size of the pre-ITQ fleet, and the trigger prices that induce entry/exit. To explore its effect on the transition and long run equilibrium, the per period fixed cost is varied from 0.000 to 0.125 across simulations. Results for these simulations are contained in Table B.8 and Figure B.4. As the fixed cost increases it becomes less profitable to operate in the fishery. This decreases the number of fishermen capable of remaining active and reduces the value of holding quota, a downward shift of the price and fleet size path. Except in the case of no fixed cost ($f = 0$), changes in the fixed cost do not appear to affect rate of downsizing. At higher fixed cost, the fleet size is smaller, more efficient, and generates less resource rents.

The time preference of fishermen affects the pricing of quota and the rate of vessel exit, but not the long run equilibrium. To examine the effect of forward-looking behavior on the transition, δ is varied between 0 and 1. The results are contained in Table B.8 and Figure B.4. Myopic fishermen ($\delta = 0$) immediately transition to the new equilibrium. As they place no value on the resale of quota nor the cost of reentry, fishermen make quota adjustments based on their fundamental valuation. As δ increases, fishermen place more weight on the value of resale and penalty of reentry, trading prices rise and the rate of exit slows.

The ex-vessel price of fish is a component of a fisherman's fundamental valuation of quota. An increase in the price of fish should produce an equivalent increase on the quota trading price, as it raises the marginal profits of all fishermen equally. Simulations reveal this pattern; see Table B.10 and Figure B.7. A 10% increase in the price for fish raises the long run trading price by 10%, and results in an upward shift in the price path. Since both the ex-vessel price and the trading price increase by an equivalent amount, there is no significant effect on the pattern of vessel exit.

Finally, setting the total allowable catch is an important consideration for fishery management. In theory, the catch target (and quota supply) is set equal to the maximum economic yield for the fishery. Selection of the supply of quota affects both the rate of fleet downsizing as well as trading prices. We vary the quota supply from 5 to 15. Table B.11 and Figure B.8 describe the transition for these corresponding values. Not surprisingly, as the quota supply increases, so does the size of the fleet in the long run equilibrium. Similarly, increasing the quota supply while holding the aggregate demand for fish constant (price of fish constant), leads to a decrease in the long run trading price for quota. This fleet is necessarily less efficient than the baseline; it supports a more cost-inefficient fleet.

3.5 Conclusions

This paper presents a dynamic transition model to study the economic fundamentals that determine the path from an initially over-capitalized fishery to the equilibrium ITQ-regime fleet structure. We characterize the fleet size, configuration, and quota trading prices that emerge during the transition. In addition, comparative dynamics exercises provide insight on the impact of key model parameters on the transition path.

We find that uncertainty over the true distribution of cost-efficiency in the population of fishermen is an important determinant for the rate of fleet downsizing. Uncertainty over relative efficiency creates uncertainty in the price of fishing quota, impacts quota purchase and sales decisions, and correspondingly the decision to remain active in the fishery. Simulations of the model find that delays in fleet restructuring occur when fishermen hold heterogeneous beliefs about their relative productivity, a likely scenario in real-world fisheries. Fleet downsizing patterns predicted by the model mirror patterns observed in some real world fisheries. The model also predicts that quota trading prices can be initially high due to the speculative premium that fishermen place on ownership of the quota asset. The speculative premium eventually disappears as fishermen learn about the true distribution of productivity through repeated interaction in the quota trading market. Other economic factors, such as entry costs, time preference, and the quota supply also affect the transition path.

Our results should be of interest to policy makers. First, the initial allocation of quota has efficiency implications during the transition. Endowing cost-inefficient fishermen with quota promotes delayed-exit strategies. This results in a slower transition to the ITQ-equilibrium, and a prolonged period during which the full efficiency benefits of ITQ management are unrealized.

In addition, concerns over social disruption and transitional costs accompanying fleet rationalization often prompt designers of ITQ programs to include trading restrictions, at least temporarily (National Research Council (1999)). This paper emphasizes the importance of learning which occurs through repeat interaction in the quota trading market. If quota trading is prohibited or restricted by regulation, learning will be delayed and the transition lengthened.

This structural model provides a baseline for further study and extensions. Calibration of the model to specific fisheries that have undergone transition to the post-ITQ equilibrium would allow an external validation of the theory in this paper and is considered a next step.

Extensions of the model should provide additional insights and refine policy recommendations. One possibility is incorporating a quota lease market. Several economists have noted that lease prices provide useful information (separate from the ownership price) that may be useful for fishery management. Anderson and Sutinen (2006) argue that restricting trading of quota to lease markets during the initial trading periods promotes price discovery and eases the transition. Grainger and Costello (2011) examine the ratio of lease to ownership prices in ITQ managed fisheries around the world. If a trading market is fully functioning, this ratio should equal the risk free interest rate. The authors note that this is not the case with most fisheries and use observed differences in the security of tradable property-rights to explain how the market adjusts for this uncertainty; for example, in U.S. fisheries an ITQ is a fishing privilege not a right. Incorporating a lease market into the dynamic transition model will allow us to explore these issues further.

A second extension would endogenize learning to the quota trading price. This adds interesting methodological novelties and has the advantage of being an internally consistent Bayesian model, which allows the formation of self-fullfilling expectations and bubbles. This also allows a deeper discussion of the effect of trading restrictions, introduced by fishery management, on the transition path and corresponding efficiency.

Bibliography

Andersen, P., Andersen, J. L., Frost, H., 2010. ITQ's in Denmark and Resource Rent Gains. Maring Resource Economics 25, 11–22.

Anderson, C. M., Sutinen, J. G., 2005. A Laboratory Assessment of Tradable Fishing Allowances. Marine Resource Economics, 1–23.

Anderson, C. M., Sutinen, J. G., 2006. The Effect of Initial Lease Periods on Price Discovery in Laboratory Tradable Fishing Allowance Markets. Journal of Economic Behavior and Organization 61, 164–180.

Anderson, L. G., 1989. Rights Based Fishing. Kluwer Academic Publishers, Ch. Conceptual Constructs for Practical ITQ Management Policies, pp. 191–209.

Anderson, L. G., 2000. The Effects of ITQ Implementation: A Dynamic Approach. Natural Resource Modeling 13 (4), 435–470.

Arnason, R., 1990a. The Icelandic Individual Transferable Quota System: A Descriptive Account. Marine Resource Economics 8, 201–218.

Arnason, R., 1990b. Minimum Information Management in Fisheries. Canadian Journal of Economics 23 (3), 630–653.

Arnason, R., 2002. Future Options for UK Fish Quota Management. Tech. rep., CEMARE Report 58.

Asche, F., Bjørndal, T., Gordon, D. V., 2009. Resource Rent in Individual Quota Fisheries. Land Economics 85 (2), 279–291.

Batstone, C. J., Sharp, B. M. H., 2003. Minimum Information Management Systems and ITQ Fisheries Management. Journal of Environmental Economics and Management 45 (2), 492–504.

Brandt, S., McEvoy, D., 2006. Distribution Effect of Property Rights: Transition in the Atlantic Herring Fishery. Marine Policy 30, 659–670.

Buck, E. H., 1995. Individual Transferable Quotas in Fishery Management. Congressional Research Services (CRS) Report for Congress, Available online: http://www.hutten.org/fw/docs/207.pdf.

Christy, F. T., 1973. Fisherman Quotas: A Tentative Suggestion for Domestic Management. Occasional Paper 19, law of Sea Institute, Honolulu, HI.

Colgan, L., Pascoe, S., 1999. Separating Resource Rents from Intra-marginal Rents in Fisheries' Economic Survey Data. Agricultural and Resource Economics Review 28 (2), 219–228.

Costello, C., Gaines, S. D., Lynham, J., 2008. Can Catch Shares Prevent Fisheries Collapse? Science 321 (5896), 1678–1681.

Deily, M. E., McKay, N. L., , Dorner, F. H., 2000. Exit and Inefficiency: The Effects of Ownership Type. The Journal of Human Resources 16 (2), 735–747.

Dixit, A., 1989. Entry and Exit Decisions Under Uncertainty. Journal of Political Economy 97 (3), 620–638.

Dupont, D. P., 2000. Individual Transferable Vessel Quotas and Efficient Restructuring of the Primary Harvesting Sector. Annals of Operations Research 94, 275–294.

Fudenberg, D., Tirole, J., 1986. A Theory of Exit in Duopoly. Econometrica 54 (4), 943–960.

Ghemawat, P., Nalebuff, B., 1985. Exit. RAND Journal of Economics 16 (2), 184–194.

Grafton, R. Q., Arnason, R., Bjørndal, T., Campbell, D., Campbell, H. F., Clark, C. W., Connor, R., Dupont, D. P., Hannesson, R., Hilborn, R., Kirkley, J. E., Kompas, T., Lane, D. E., Munro, G. R., Pascoe, S., Squires, D., Steinshamn, S. I., Turris, B. R., Weninger, Q.,

2006. Incentive-Based Sustainable Fishery Management. Canadian Journal of Fisheries and Aquatic Science (63), 699–710.

Grafton, R. Q., Squires, D., Fox, K. L., 2000. Private Property and Economic Efficiency: A Study of a common Pool Resource. Journal of Law and Economics 43, 671–714.

Grainger, C., Costello, C., 2011. The Value of Secure Property Rights: Evidence from Global Fisheries. NBER Working Paper No. 17019.

Harrison, M., Kreps, D., 1978. Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations. The Quarterly Journal of Economics 92 (2), 323–336.

Larkin, S., Milon, J., 2000. Tradable Effort Permits: A Case Study of the Florida Spiny Lobster Trap Certificate Program. Proceedings of the 2000 International Institute of Fisheries Economics and Trade, Corvallis, available online: www.oregonstate.edu/dept/IIFET/2000/papers/larkin.pdf.

Lian, C., Singh, R., Weninger, Q., 2009. Fleet Restructuring, Rent Generation and the Design of Individual Fishing Quota Programs: Empirical Evidence from the Pacific Coast Groundfish Fishery. Marine Resource Economics 24 (4), 329–359.

Lindner, R. K., Campbell, H., Bevin, G., 1992. Rent Generation During the Transition to a Managed Fishery: The Case of the New Zealand ITQ System. Marine Resource Economics 7, 229–248.

Matulich, S., Mittelhammer, R., Reberte, C., 1996. Toward a More Complete Model of Individual Transferable Fishing Quotas: Implications of Incorporating the Processing Sector. Journal of Environmental Economics and Management 31, 112–128.

Montgomery, W. D., 1972. Markets in Licenses and Efficient Pollution Control Programs. Journal of Economic Theory 5 (3), 395–418.

Morris, S., 1996. Speculative Investor Behavior and Learning. The Quarterly Journal of Economics, 1111–1133.

Moxnes, E., 2007. Individual Transferable Quotas Versus Auctioned Seasonal Quotas: An Experimental Investigation. Economic Science Association, World Meeting, Rome, available online: www.ifi.uib.no/sd/.

Munro, G. R., Scott, A. D., 1985. Handbook of Natural Resources and Energy Economics. Ch. The Economics of Fisheries Management.

National Oceanic and Atmospheric Administration, 2010. NOAA Catch Share Policy. Available online: www.nmfs.noaa.gov/sfa/domes_fish/catchshare/docs/noaa_cs_policy.pdf.

National Oceanic and Atmospheric Administration, 2011. Catch shares. Available online: www.nmfs.noaa.gov/sfa/domes_fish/catchshare/index.htm.

National Research Council, 1999. Sharing the Fish: Toward a National Policy on Individual Fishing Quotas. National Academy Press.

Newell, R. G., Papps, K. L., Sanchirico, J. N., 2007. Asset Priceing in Created Markets. American Journal of Agricultural Economics 89 (2), 259–272.

Newell, R. G., Sanchirico, J. N., Kerr, S., 2002. Fishing Quota Markets. Journal of Environmental Economics and Management 49 (3), 437–462.

Squires, D., Kirkley, J., Tisdell, C. A., 1995. Individual Transferable Quota as a Fisheries Management Tool. Review in Fisheries Science 2 (2), 141–169.

Vestergaard, N., Jensen, F., Jørgensen, H. P., 2005. Sunk Cost and Entry-Exit Decisions Under Individual Transferable Quotas: Why Industry Restructuring is Delayed. Land Economics 81 (3), 363–378.

Weninger, Q., 1998. Assessing Efficiency Gains from Individual Transferable Quotas: An Application to the Mid-Atlantic Surf Clam and Ocean Quahog Fishery. American Journal of Agricultural Economics 80, 750–764.

Weninger, Q., Just, R. E., 1997. An Analysis of Transition from Limited Entry to Transferable Quota: Non-Marshallian Principles for Fisheries Management. Natural Resource Modeling 10 (1), 53–83.

Weninger, Q., Just, R. E., 2002. Firm Dynamics with Tradable Output Permits. American Journal of Agricultural Economics 84 (3), 572–584.

Weninger, Q., Waters, J. R., 2003. Economic Benefits of Management Reform in the Northern Gulf of Mexico Reef Fish Fishery. Journal of Environmental Economics and Management 46, 207–230.

CHAPTER 4. MODEL VALIDATION OF THE PURE-CHARACTERISTICS VERTICAL SORTING MODEL

4.1 Introduction

There are well documented concerns with using the hedonic property value model to measure the welfare effect from a large or discrete change (see Freeman (2003)). Potential issues range from identification of the second stage bid function to violations in assumed continuity of variables included in the hedonic price schedule. These concerns have raised new interest in alternative methods for estimating welfare. Equilibrium sorting models have been suggested as an alternative that may avoid many of these issues.

By introducing additional structure, equilibrium sorting models estimate preferences in a general equilibrium framework (Sieg et al. (2004)). Preference estimates along with the added structure are sufficient to predict equilibria that may emerge from a non-marginal change; such as the air quality changes envisioned by the 1990 amendments to the Clean Air Act (Smith et al. (2004)). These computed equilibria provide information about how prices and choices are altered in response to an exogenous change. This allows the economist to recover a general equilibrium estimate of welfare using standard Hicksian measures of compensating or equivalent variation. The accuracy of general equilibrium welfare measurement depends, of course, on how well the model recovers preferences and the accuracy of the prediction from the computed equilibria. However, while these models are capable of predicting how people and markets will adjust to large-scale policy change, we have no evidence on the accuracy of their predictions. Thus, before these models can be used to make recommendations in substantive policy debate, it is important to examine the predictive performance of this class of models. The goal of this work is to provide the first effort and predict performance.

More generally, model validation is important. Keane (2010) argues that economists should increase efforts in validating structural econometric models, as this is an important step in gaining general use and acceptance. This paper is the first attempt at external model validation of the pure-characteristics vertical sorting model established by Epple and Sieg (1999). Using the cleanup of Luke Air Force Base (a deleted Superfund site in Maricopa County, AZ) as a quasi-experiment, data surrounding cleanup of this site is used to compare the observed and predicted sorting of households, prices, housing expenditures, and income. In-sample tests are also performed to see how well the data used in the estimation of model parameters fits the structure of the model. Welfare analysis provides an additional source of model validation, and highlights a proposed strength of the sorting model, namely its ability to analyze the distributional and spatial effects of government policy.

In this paper, data surrounding remediation of the Superfund site is used to estimate three models. Models I and II use data from before and after the cleanup, respectively, to estimate the model parameters and predict the equilibrium in the other state of the world. For example, Model I parameter estimates are used to perform a prospective analysis of cleanup of Luke Air Force Base. Model III stacks data from before and after the change to perform a joint estimation of the model, which allows unobservables to vary over time and provides a robustness check for the parameter estimates in Models I and II.

The results suggest that the ability of the sorting model to predict the response of households to large-scale changes does not come without a cost. In-sample and out-of-sample results find an average bias (absolute deviation) of nearly 25% for income and 20% for housing expenditures. Predicted prices also miss the housing bubble observed in the early 2000s. The model does appear to capture the movement of population shares in response to the cleanup. The mobility of households is sensitive to parameter estimates. Model II predicts much more movement by households than Model I. The spatial distribution of general equilibrium welfare is a function of the structure of the model. Estimates of community price indices define the substitutes for households. Mean general equilibrium welfare values vary considerably across Maricopa County; from −\$52.78 to \$24.50 for the remediation of soil and groundwater contaminants at Luke Air Force Base. Averaged across the community populations, the mean general equilibrium willingness-to-pay was slightly higher than the partial equilibrium counterpart, \$0.19 and \$0.01 respectively.

The remainder of the paper is outlined as follows. The next section provides a discussion of the background literature on the vertical sorting model and discusses the setting of the model validation exercise. Section 4.3 presents the structure of the pure-characteristics vertical sorting model. In Section 4.4 the data and setting for the model validation exercise is described. Results from model estimation, in-sample and out-of-sample tests, and welfare analysis are discussed in Section 4.5. Section 4.6 discusses the results and next steps for research.

4.2 Background and Previous Literature

Equilibrium sorting models are a class of econometric models, which are based on the idea that households "vote with their feet" (Tiebout (1956)). That is, households make location choices, not just on the structural characteristics of a home, but also on local neighborhood characteristics, taxes, environmental amenities, etc (i.e. nonmarket attributes). As such, household location choices reveal information about preferences. This class of models uses the principle of market equilibrium and information on household location choices, prices, and community-specific attributes to elicit the preferences of heterogenous households. Equilibrium sorting models generally fall into one of two categories, the pure-characteristics "vertical" sorting model and the discrete-choice "horizontal" sorting model.

The pure-characteristics vertical sorting model is a structural econometric model based on a discrete-continuous choice framework. In this model, households select their preferred community and amount of housing services to consume. In the vertical sorting model, households objectively agree on the rank ordering of communities from lowest to highest "quality," which is similar to the concept of vertical product differentiation in the I.O. literature. Epple and Sieg (1999) provide the first estimation of this model using a two-stage estimator. The main contribution of their paper is the recognition that the structure of the theoretic model is sufficient to provide point identification of parameter estimates. Sieg et al. (2004) propose a single-stage estimator using simulated GMM that incorporates spatial information on housing expenditures. The authors emphasize that the structure of the model allows identification of a general equilibrium Hicksian measure of willingness-to-pay. Several additional methodological advances have occurred since this paper (see Klaiber and Smith (2010), Kuminoff (2009), among others). For example, Walsh (2007) demonstrates that the provision of the public good can be modeled as endogenous to the sorting of households. He finds that a policy to preserve open-space can actually lead to a net decrease as households substitute away from the private provision to public provision.

The discrete-choice horizontal sorting model, established by Bayer et al. (2004), is based on the discrete-choice framework. The econometrician estimates the probability that households select a "type" of housing unit. Because of household-specific idiosyncratic unobservables, each household holds a subjective rank ordering of communities by "quality." This is similar to the concept of horizontal product differentiation in the I.O. literature. The horizontal sorting model operates under a probabilistic concept of equilibrium, where markets clear in expectation. As such, the model predicts the probability of household movement in response to a policy change. There are several empirical applications of this model being used to elicit the value of environmental attributes, such as open-space and air quality (Klaiber and Phaneuf (2010), Tra (2010), Bayer et al. (2009)).

The literature has yet to perform model validation for the pure-characteristics vertical sorting model. Banzhaf and Walsh (2008) provide a partial exception. The authors use a reduced form model to test the direction of changes in population and income implied by the structure of the sorting model. This, however, does not provide direct insight into the predictive performance or the sorting mechanism of the model. Our paper uniquely exploits the full structural model (in estimation and prediction) and data from a quasi-experiment to test how well the model performs at inside sample and outside sample validation exercises.

Estimation of the vertical sorting model imposes structure on the data. Specifically, identification of model parameters requires an assumption on the joint distribution that describes the heterogeneity of taste and income, and thereby housing expenditures. Inside sample validation examines how well the data fits this distributional assumption. The outside sample validation examines how well the sorting mechanism predicts the observed movement of households and prices. The comparison provides evidence on the predictive power of the sorting mechanism and its effect on general equilibrium welfare estimates. The cleanup of Luke Air Force Base (a deleted Superfund site) in late 2000 offers a useful quasi-experiment.

Luke Air Force Base, located in Glendale, AZ (Dysart Unified District), provides advanced flight training for U.S. Air Force fighter pilots. Up until late 2000, this site was a Superfund site and posed a serious health risk. According to the U.S. Environmental Protection Agency (2011b), discharges and waste disposal from normal operations created soil and possible groundwater contamination. Waste oils and volatile organic compounds, including radiological waste were discovered at the site. About 8,000 people were in direct contact with the base and faced potential exposure from soil contamination. Another 1.5 million people were potentially affected through groundwater contamination. The Phoenix groundwater basin, which runs below Luke Air Force Base, supplies water to Phoenix, Goodyear, and Youngtown, AZ. Site cleanup began in September 27, 1990 and was officially completed September 25, 2000.

The Superfund program is an alias for the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA) that was enacted by Congress in 1980. This Act charged the Environmental Protection Agency (EPA) with three tasks. First, establish a process for identifying hazardous waste sites that present a serious health risk to humans and/or the ecosystem. Second, provide liability for persons responsible for release of hazardous waste material. And third, create a trust (a Superfund) to fund remediation of hazardous waste sites when no such persons can be assigned liability.[1]

The EPA developed the Hazard Ranking System (HRS) to identify sites that should be added to the National Priority List (NPL), a list of waste sites that are given priority for cleanup. The HRS score is an index, ranging from 0 to 100, designed to account for up to four pollution pathways (groundwater migration, surface water migration, air migration, and soil exposure).[2] The score is based on a series of factors including the likelihood of exposure, the potential health risks associated with the hazardous waste, and the ecosystem and/or group of

[1] Funding for this trust comes from a tax on chemical and petroleum industries. $1.6 billion was collected in the first five years (U.S. Environmental Protection Agency (2011b)).

[2] In 1986, the Superfund Amendments and Reauthorization Act (SARA) required the EPA to revise the HRS to "ensure that it accurately assessed the relative degree of risk to human health and the environment posed by uncontrolled hazardous waste sites that may be place on the NPL." (U.S. Environmental Protection Agency (2011b)) Guidelines for the original and revised HRS are published in the Federal Register (July 16, 1982, 47 FR 31180 and December 14, 1990, 55 FR 51532).

people facing exposure. Based on this ranking system, any site that receives a score of 28.5 or higher is eligible for placement on the NPL.

4.3 The Pure-Characteristics Vertical Sorting Model

The following is a description of the pure-characteristics vertical sorting model, as outlined in Epple and Sieg (1999) and Sieg et al. (2004).

4.3.1 Structure of the Model

Consider an economy that consists of a continuum of households living in a Metropolis. The Metropolis contains a *heterogeneous* landscape that is separated into J communities with fixed boundaries that may vary in size. Communities $j = 1, \ldots, J$ differ in their provision of a composite public good, g_j, and price for *homogeneous* units of housing, p_j. Epple and Sieg (1999) define the composite public good as a combination of "locally provided public goods, environmental amenities, and any other community-specific attributes" (pp. 649). In the following, the term public good is used to refer to this composite, which will be measured as an index.[3]

Households have preferences over the public good, housing h, and a composite private good whose price is normalized to 1. Households differ in their endowment of income and taste for the public good, α. The heterogeneity of households is described by the joint distribution of y and α, which is assumed to have the continuous density, $f(\alpha, y)$.

Households face a decision at both the extensive and intensive margin. At the extensive margin households select the community that provides the preferred level of the public good. Conditional on community choice, a household selects the quantity of housing to consume. Housing can be consumed in any nonnegative amount. All remaining income is spent on consumption of the numeraire.

Let $V(\alpha, y, p_j, g_j)$ denote the indirect utility household (α, y) receives from selecting community j with price-public good pair (p_j, g_j) and consuming an optimal amount of housing

[3]Households agree on both the components and weights of the public good index. As such, households also agree on the rank ordering of communities. Berry and Pakes (2007) label this property the "pure characteristics" approach to model demand for a differentiated product.

and the numeraire. Let $\mathfrak{J}$ denote the set of communities. Then the household's problem is equivalent to selecting the community $j \in \mathfrak{J}$ that maximizes indirect utility.

$$j^* = \arg\max_{j \in \mathfrak{J}} V\left(\alpha, y, p_j, g_j\right) \tag{4.1}$$

To close the model, household choices must be consistent with market clearing in the housing market for each community. Let $H_j^D(p_j)$ denote the aggregate demand for housing in community j at price p_j. Given the distribution of heterogeneity of households, this is simply,

$$H_j^D(p_j) = \int_{C_j} h\left(p_j, y\right) f\left(\alpha, y\right) d\alpha dy / P(C_j)$$

where C_j denotes the set of households that choose community j and $P(C_j)$ the corresponding mass of households.[4]

Let $H_j^S(p_j)$ denote the aggregate supply of housing in community j at price p_j. Then the housing market in community j will clear if and only if,

$$H_j^S(p_j) = H_j^D(p_j) \tag{4.2}$$

Definition 2. *The model is in a* **sorting equilibrium** *if, given market prices and public good levels, every household is in their preferred community, consuming their optimal quantity of housing and the numeraire given their budget constraint, and all housing markets clear; that is equations (4.1) and (4.2) are satisfied for all households and all communities.*

4.3.2 The Single-Crossing Condition

The single-crossing condition is a necessary condition for the existence of a sorting equilibrium and is used in the estimation of model parameters. Kuminoff et al. (2010) provide a detailed discussion of the single-crossing condition within the sorting literature. Ellickson (1971) was the first to derive restrictions on preferences that would support such an equilibrium. Given that households vary only in their endowment of income, indirect indifference curves in the price-public good space must be monotonically increasing in income and thereby only cross once. Westoff (1977) provides a formal proof for the existence of a sorting equilibrium given households differ in income. Early extensions of the model all maintain this

[4] $P(C_j) = \int_{C_j} f(\alpha, y) d\alpha dy$.

single source of heterogeneity (Epple et al. (1994), Epple et al. (1993), and Epple and Romer (1991)). This formulation results in equilibria with perfect stratification of households by income. Epple and Platt (1998) generalized the model to include heterogeneity of preferences and income, which results in a more realistic description of household sorting; stratification conditional on tastes. Equilibrium, in this model, is characterized by a generalized version of the single-crossing condition.

Definition 3. *The* **single-crossing condition** *requires that indirect indifference curves in price-public good space cross only once. With heterogeneity in preferences for the public good and income, this is equivalent to a restriction on the conditional slope of the indirect indifference curve. That is, conditional on taste, indirect indifference curves are monotonically increasing in income, and conditional on income, are monotonically increasing in taste.*

To date, the literature has not provided a formal proof of existence of a sorting equilibrium when households have heterogenous tastes and income. Despite this, the literature provides several numeric examples where an equilibrium is found; see Klaiber and Smith (2010), Kuminoff (2009), Sieg et al. (2004), Smith et al. (2004), Epple et al. (2001), and Epple and Sieg (1999) among others.

The single-crossing condition can be given an economic interpretation. Equation (4.3) represents the slope of an indirect indifference curve in price-public good space.

$$\left.\frac{\partial p_j}{\partial g_j}\right|_{V=\bar{V}} = -\frac{\partial V/\partial g_j}{\partial V/\partial p_j} = \frac{1}{h(p_j, y)}\left[\frac{\partial V/\partial g_j}{\partial V/\partial y}\right] \tag{4.3}$$

Appealing to Roy's identity, this slope can be decomposed into two pieces, the inverse housing demand and Marshallian virtual price for the public good. As stated in Kuminoff et al. (2010) the single-crossing condition implies that "the Marshallian virtual price [for the public good], per unit of housing, is strictly increasing in income and in preference for public goods relative to private goods" (pp. 16, comment in brackets added).

Kuminoff et al. (2010) also relate the single-crossing condition to the "Willig (1978) condition that is often applied together with weak complementarity to identify Hicksian willingness-to-pay for changes in the public good" (pp. 16).[5] The Willig condition requires the slope of indirect indifference curves in price-public good space to be independent of income. This implies that the Marshallian virtual price for the public good, per unit of housing, is also independent of income. That is, marginal increases in income that raise the virtual price must be fully offset by increases in housing demand so that the per unit value is unaltered. The single-crossing condition requires that housing demand is less responsive than the virtual price to changes in income, so that this slope is strictly increasing in income.[6]

As outlined in Epple and Platt (1998), the single-crossing condition implies that in equilibrium three properties characterize the sorting of households: (1) boundary indifference, (2) stratification, and (3) ascending bundles.

- Boundary indifference states that there exist households, (α, y) pairs, that are indifferent between communities that are adjacent in their relative housing prices and provision of the public good.

- Stratification asserts that households sort across the J communities such that they are stratified by income given taste and taste given income.

- The ascending bundles property requires that the price for housing services and the provision of the public good move together. That is, ranking communities by price is equivalent to ranking the communities by quality (provision of the public good).

These three properties are necessary for the existence of a sorting equilibrium and are used in the estimation of the vertical sorting model.

4.3.3 Model Parameterization

To estimate the model, we make use of the parameterization of indirect utility used in Epple and Sieg (1999), Sieg et al. (2004), Smith et al. (2004), among others.

$$V\left(\alpha, g, p, y\right) = \left\{\alpha g^{\rho} + \left[\exp\left(\frac{y^{1-\nu}-1}{1-\nu}\right)\exp\left(-\frac{\beta p^{\eta+1}-1}{1+\eta}\right)\right]^{\rho}\right\}^{\frac{1}{\rho}} \tag{4.4}$$

[5] Smith and Banzhaf (2004) provide a helpful discussion on the interpretation of the Willig condition, weak complementarity, and the recovery of Hicksian willingness-to-pay.

[6] Smith and Banzhaf (2007) provide an alternative interpretation based on the expansion of indifference curves as income changes.

This is a standard CES function where ρ denotes a measure of substitution between housing and public good, η and ν are the constant price and income elasticities of housing demand, and β is a scaling parameter for housing demand. These parameters are common across all households.

This specification has some helpful properties. First, it separates the utility effect of private and public goods. The first term in the CES function represents the utility from locally provided public goods. The second term is the contribution of private goods (housing and the numeraire) to utility. Due to this separability the CES function generates simple housing demand equations that are independent of community-specific amenities. Following Roy's identity, housing demand is simply,[7]

$$h(p_j, y) = \beta p_j^{\eta} y^{\nu} \tag{4.5}$$

Second, the CES function satisfies the single-crossing condition if $\rho < 0$. Sieg et al. (2004) and Kuminoff et al. (2010) note that this condition offers the opportunity to test the underlying theory.

Two additional parameterizations are needed to estimate the model, the distribution describing the heterogeneity of households and the public goods index. Universally the literature uses the bivariate log normal distribution to describe the heterogeneity of households. The public good is approximated by a linear function of its components.

4.3.4 Sorting

To understand the strengths and weaknesses of the vertical sorting model, it is useful to discuss how the structure of the model induces sorting. Boundary indifference states that there exists households that are indifferent between adjacent (by quality/price ranking) communities. For a household to be indifferent between communities j and $j+1$, it must be that,

$$V(\alpha, y, g_j, p_j) = V(\alpha, y, g_{j+1}, p_{j+1})$$

[7] In the estimation of the model, relative communities prices are used in the housing demand equation. The parameter β rescales to correct for this additional normalization.

Using the specification of indirect utility in equation (4.4), boundary indifference can be expressed to separate the household parameters from the community characteristics,

$$M(\alpha, y) \equiv \underbrace{\ln(\alpha) - \rho\left(\frac{y^{1-\nu}-1}{1-\nu}\right)}_{\text{Household parameters}} = \underbrace{\ln\left(\frac{Q_{j+1}-Q_j}{g_j^{\rho}-g_{j+1}^{\rho}}\right)}_{\text{Community characteristics}} \equiv K_j \tag{4.6}$$

where $Q_j = \exp\left(-\rho(\beta p_j^{\eta+1}-1)/(\eta+1)\right)$ for $j = 1, \ldots, J-1$. The left-hand side of equation (4.6) represents the loci of households that are indifferent between communities j and $j+1$. The term $M(\alpha, y)$ captures this combination of preferences and income, and thereby acts as a household preference index. The right-hand side only depends on the characteristics of the communities j and $j+1$. The term K_j denotes an index for these community characteristics and is used to define a cut point between the communities j and $j+1$.[8] If $M(\alpha, y)$ exceeds K_j, the household will prefer community $j+1$ to community j.

A household with a higher $M(\cdot)$ will prefer a community with a larger provision of the public good, despite the higher community price. Given the household preference index and cut points we can easily sort households. Sorting is depicted in Figure C.1. The bold lines denote the combinations of taste and income under which a household is indifferent between two communities. Community j will consist of all households such that $M(\alpha, y)$ is contained between the boundaries K_{j-1} and K_j; $C_j = \{(\alpha, y)|K_{j-1} \leq M(\alpha, y) \leq K_j\}$. This sorting captures a household's trade-off between consuming a large quantity of housing and/or the numeraire (at a low price) but receiving a low quality public good, and consuming a high quality public good but facing a high per unit price of housing.

Figure C.1 illustrates some interesting sorting behavior based on household preferences. Consider three households, a, b, and c that sort between three communities, as shown in the figure. Household a has the smallest endowment of income. Despite this fact, it selects the most expensive community. This is due to the household's strong taste for the public good (it has the strongest preference of the three households). Household c, on the other hand, has the highest income yet selects the lowest quality community. This imperfect stratification of households by income/preferences is an important feature of the sorting model.

[8]To complete sorting, the following end points are added, $K_0 = -\infty$ and $K_J = \infty$.

4.3.5 Estimation

Following Sieg et al. (2004) and Smith et al. (2004), estimation of the pure-characteristics vertical sorting model is based on simulated GMM. This method stacks three sets of moment conditions derived from the theoretic model and then minimizes the distance between the empirical and fitted moments. Estimation of the model parameters relies on fitting (1) a linear approximation for the public goods index, (2) the distribution of income across communities, and (3) the distribution of housing expenditures. The following is a brief description of each of these moment conditions.

The first set of moment conditions minimize the distance between an "implied" public goods index, $\tilde{g}_j$, and a linear approximation.

$$\tilde{g}_j - \gamma X_j = \xi_j \tag{4.7}$$

where X_j denotes the observable community-specific amenities and ξ_j the characteristics of community j that are observable to households but not the econometrician. ξ_j is assumed uncorrelated with X_j and small enough that it does not affect the rank ordering of communities.

In Epple and Sieg (1999), the authors recognize despite that the provision of public good is only partially observable, the structure of the model still makes estimation possible. Observed data on household location choices, prices, community-specific attributes, and population shares can be used to calculate the level of public goods implied by the data. Given an estimate for the cut points, K_j, and the parameters of the model, the boundary indifference property is used to solve for the implied public goods.[9]

$$\tilde{g}_j = \left\{ \tilde{g}_1^{\rho} - \sum_{i=2}^{j} (Q_i - Q_{i-1}) \exp(-K_{i-1}) \right\}^{\frac{1}{\rho}} \qquad \text{for } j = 2, \ \ldots, \ J \tag{4.8}$$

This function defines an analytic recursive algorithm that is a function of the public good in the previous community. As such, the public goods index in the lowest ranked community, g_1, is also a parameter of the model.

[9] This is the same method used in Klaiber and Smith (2010) and Kuminoff (2009). Alternatively, Sieg et al. (2004) suggest inverting the population share equations, which generate implicit (recursive) functions of the implied public good that exactly match population shares.

The second set of moment conditions minimize the difference between empirical and fitted log income percentiles.

$$\ln y_j^E(q) - \ln y_j^F(q) = 0 \tag{4.9}$$

where the superscripts E and F denote empirical and fitted information and q denotes the percentile of interest. The literature focuses on the use of the 25th, 50th, and 75th percentiles for these moments, though more can be included in the estimation depending on the quality of the data.

Finally, the third set of moment conditions minimize the difference between empirical and fitted log housing expenditure percentiles. While empirical housing expenditures are derived from the transaction prices of homes in a community, fitted housing expenditures are the product of the aggregate quantity of housing consumed and the per unit price for housing in a community. From our specification of indirect utility, we make use of the housing function in equation (4.5), which gives us,

$$\ln E_j^E(q) - \ln \beta - (\eta + 1) \ln p_j - \nu \ln y_j^F(q) = 0, \tag{4.10}$$

where $E_j^E(q)$ denotes the qth empirical quantile for housing expenditures in community j. Again the focus is on the 25th, 50th, and 75th percentiles for this set of moments. Note that the qth fitted expenditure on housing uses the qth fitted log income.

Estimation of the model also requires the use of instrumental variables. Instruments are needed for a couple of reasons. First, a household's choice of community is equivalent to the choice of housing price; as is apparent from equation (4.1). This makes the price-quantity decision a joint choice. Second, there are omitted variables such that the unobservable component of the public good is correlated with other variables.

Sieg et al. (2004) and Smith et al. (2004) recommend the use of polynomial functions on the price ranking of communities. Given the ascending bundles property, the ranking of a community by quality (unobservable) is equivalent to its ranking by price (observable). This uses the equilibrium properties of the model to generate instruments. Following Kuminoff (2009), Chebychev polynomials on the price ranking of communities are used to generate instruments. These polynomials have the benefit that the instruments will be strongly correlated with price but only weakly correlated with each other, ensuring that each additional instrument provides new information to the estimation.

4.3.6 Computing Counterfactual Equilibria

Sieg et al. (2004) outline a method for computing counterfactual equilibria. This method is separated into two steps. First, generate a baseline simulated landscape of household sorting based on observed data and estimated model parameters. Second, given an exogenous change in public goods, re-sort households and solve for the new market clearing set of prices. The following describes each step.

To generate the baseline, a large sample of simulated households is drawn from $f(\alpha, y)$ using the parameter estimates.[10] Given the baseline level of prices, implied public good provision, and population shares, these simulated households are sorted across the Metropolis as outlined in Section 4.3.4.

An important consideration in computing the counterfactual equilibrium is the parameterization for the aggregate housing supply function. Consider the following parameterization used in Sieg et al. (2004) and Epple and Platt (1998),[11]

$$H_j^s(p_j) = l_j p_j^{\tau}$$

where τ denotes the price elasticity of housing supply and l_j a supply scaler. The baseline supply of housing is calibrated to the baseline demand (the initial equilibrium).

Given a perturbation to the provision of the public good, the hierarchial structure of the vertical sorting model allows the solution to this complex system of nonlinear equations to be reduced to a search over a single parameter. The solution method is simple. Given an initial guess for the price in the lowest (price/quality) community, market prices for the next $J-1$ markets are solved using the recursive structure of the model. This reduces the search problem to iterating the price in the lowest community until the Jth market clears.

[10] In this paper we use four million draws from the joint density of y and α.

[11] Epple et al. (2001) parameterize the aggregate housing supply function using a constant returns to scale Cobb-Douglas specification.

4.3.7 Welfare Analysis

One of the proposed strengths of the sorting model is the ability to analyze the full distributional and spatial effects of government policy. Sieg et al. (2004) emphasize that the structure of the vertical sorting model allows recovery of a general equilibrium measure of Hicksian willingness-to-pay.

To illustrate, consider an exogenous improvement in the quality of community j, from g_j to $\widetilde{g}_j$. In the partial equilibrium setting, Hicksian willingness-to-pay is defined as,

$$V(\alpha, y - WTP_{PE}, \widetilde{g}_j, p_j) = V(\alpha, y, g_j, p_j),$$

which restricts households to their original community choice, j, and housing prices.

In a general equilibrium setting, households respond to the quality change through (possibly) new community choices. The change in the public good and movement of households cause housing markets to adjust and relative community prices to change. The proposed general equilibrium measure includes both of these effects,

$$V(\alpha, y - WTP_{GE}, \widetilde{g}_k, \widetilde{p}_k) = V(\alpha, y, g_j, p_j),$$

where k denotes the community choice in the new equilibrium.

Characterizing this general equilibrium welfare measure requires the solution to the counterfactual equilibrium, e.g. the re-sorting of households and new equilibrium prices. As outlined in Section 4.3.6, this uses simulated households, (α, y) pairs, drawn from the estimated joint distribution. For each simulated household, the partial and general equilibrium WTP is simple to calculate.[12] Once WTP is calculated for each simulated household, this information can be analyzed along any dimension of interest, such as the mean WTP by community.

[12] The CES specification of indirect utility generates an analytic expression for WTP.

$$WTP_{GE} = y - \left\{ \frac{1-\nu}{\rho} \ln(V_0^{\rho} - \alpha \widetilde{g}_k^{\rho}) - \frac{1-\nu}{1+\eta} (\beta \widetilde{p}_k^{\eta+1} - 1) + 1 \right\}^{\frac{1}{1-\nu}}$$

4.4 Data and Setting

The quasi-experiment focuses on Maricopa County, which is located in south-central Arizona. This county is the fourth most populated in the U.S. and contains eight cities with populations in excess of 100,000. Maricopa County also contains 56% of Arizona's Superfund sites, seven sites in total. Currently these sites are in various stages of remediation. 57.1% have been classified as completed and/or deleted from the National Priority List. The remaining 42.9% are still undergoing remediation. Table C.1 contains a history of remediation for Superfund sites in Maricopa County and their corresponding HRS scores.

Movement from one stage of remediation to the next provides information to households about the local provision of environmental amenities. Households react to this information in the housing market through (possibly) new location choices. In a general equilibrium setting, this leads to the formation of a new sorting of households and new equilibrium prices. Then, assuming that incidental changes in other public good attributes do not induce a re-sorting of households, the dates in the table allow the time-horizon to be partitioned into a series of equilibria.

The inside and outside sample tests in this paper focus on the equilibria surrounding the cleanup of Luke Air Force Base, located in Glendale, AZ. Using the dates from Table C.1, data for these equilibria is collected over two time periods; the initial equilibrium (02/17/1998 to 09/25/2000) and the new equilibrium (09/25/2000 to 05/01/2003). In the following, Model I refers to the estimation of the model using data from the initial equilibrium and Model II, data from the new equilibrium. A joint estimation of the sorting model, using stacked data from before and after cleanup, is included as a robustness check for the parameter estimates and is labeled Model III.

The school district is treated as the unit of economic analysis. That is, household location decisions correspond to a choice of school district. To estimate the model, we collect data on household characteristics by school district, quality measures for local public education and air quality, Superfund site characteristics, and data on housing markets.

4.4.1 The Community

The definition of community is an important modeling consideration. First, the community is the smallest unit of economic analysis in the model. Second, its definition can affect both parameter estimates and the nature of sorting. Kuminoff (2009) finds that the definition of the choice set has strong consequences on the sorting of households. Most applications use the school district as the choice of community (Kuminoff (2009), Sieg et al. (2004), Smith et al. (2004), among others), though alternative choices exist in the literature. Banzhaf and Walsh (2008), for example, create artificial communities using circles with one-mile and half-mile diameters, evenly spread over urban areas in California (6,218 and 25,166 communities respectively). Walsh (2007) constructs 91 neighborhoods in Wake County, North Carolina from information on jurisdictional boundaries, school attendance zones and major roadways. Epple and Sieg (1999) define communities as the cities and townships in the Boston Metropolis Area (92 communities in total).

In this paper, unified, secondary, and elementary school districts are combined to define the community choice set. This provides a large coverage of Maricopa County, excluding invalid areas such as Indian Reservations and uninhabitable desert regions. In addition to these excluded areas, two school districts were dropped from the choice set because no housing transactions were observed in the database between 1998 and 2003. This left a choice set of 24 communities, as depicted in Figure C.2.[13] Throughout the remainder of the paper, the terms community and school district are used interchangeably.

4.4.2 Housing Markets and Prices

Information on housing markets in Maricopa County was constructed using a nearly exhaustive database of housing transactions for single-family homes between 1974 and 2004. The database was assembled from three sources. The main source, DataQuick, is a private firm that collects and maintains records on housing transactions for sale to realtors, financial and academic institutions. DataQuick's records contain information on housing characteristics and transaction prices. This data was supplemented with additional information collected from

[13] Klaiber and Smith (2010) focus on elementary school districts in Maricopa County (46 communities).

County Assessor and Parcel records. The housing markets database is used to calculate the empirical distribution of annualized housing expenditures and estimate price indices that represent community prices in the model estimation. The following describes this process.

Households are modeled as renters of housing units. As such, transaction prices must be converted to imputed rents. Following Porterba (1992) the ownership price of housing is scaled according to a "user cost,"

$$r_n = \underbrace{\left[i + \tau_p + \beta + \delta + m - \pi\right]}_{\text{user cost}} \times p_n^H$$

where p_n^H denotes the transaction price for housing unit n and r_n the same units corresponding rental price. The user cost is a linear combination of the interest rate i, property taxes τ_p, risk preference β, depreciation on the housing asset δ, the maintenance rate m, and inflation π. As formulated, a higher interest rate raises the user cost, while inflation lowers it. To calculate the user costs for Models I and II, data was collected on the 30-year average FHA mortgage rate, the average property tax rate in Maricopa County, and the CPI for housing in urban areas of the Western U.S. Values for the remaining parameters are drawn from Porterba (1992). The corresponding values provide a user cost of 13.6% for Model I and 12.8% for Model II. Table C.2 details the values used in the calculation of user costs.

Imputed rents are equivalent to annualized housing expenditures. Separating these rents by school district, calculation of the empirical distribution of housing expenditures by school district is straightforward.

The vertical sorting model operates under the assumption that housing is homogeneous within and across communities. Under this assumption, deviation in prices across the Metropolis are defined by differences in the provision of the public good across school districts. The average rent from the transactions data will not be an appropriate index for this relationship as it does not control for structural differences in housing. Consider Table C.3, which contains summary statistics on the housing characteristics of five school districts in the household's choice set. The mean and median housing characteristics vary considerably across these school districts. For example, Scottsdale Unified District tends to have larger houses on smaller lots than Aguila Elementary District, and has over triple the expenditures on housing.

Sieg et al. (2002) propose the use of a fixed effects hedonic regression to estimate price indices that represent community prices.

$$\ln(r_{jtn}) = \beta' S_{jtn} + \sum_j \delta_j C_{jn} + \sum_t \delta_t Y_{tn} + \epsilon_{jtn}$$

In the above, S_{jtn} denotes the structural characteristics (bathrooms, liveable space, etc.) of house n in community j sold in year t. C_{jn} is a dummy variable that equals 1 if housing n is sold in community j and 0 otherwise (community fixed effects) and Y_{tn} is a dummy variable that equals 1 if housing n is sold in year t and 0 otherwise (time fixed effects). Lastly ϵ_{jtn} is an idiosyncratic error component capturing market unobservables. The community fixed effects, δ_j's, represent the community price indices.[14]

Structural control variables include lot acres, square footage, floors, bathrooms, age, garage, and pool. Table C.4 contains the results for the fixed effects hedonic regressions. Most parameter estimates are significant and all have the expected sign. In Model II, three community fixed effects have small t-stats (less than 2), which is partially explained by an unusually small number of observations for these school districts in Model II. As depicted in Table C.5, these three school districts experience an estimated two to three position change in rank ordering. For example, Higley Elementary School District moved from position 14 to 11, a three position loss in relative price ranking. All three of these communities fall within the middle of the household's choice set.

One natural concern is that anticipation of cleanup of a Superfund will affect the market price for housing and bias community price estimates.[15] To test the robustness of these results an additional set of fixed effects regressions were estimated, dropping data 6 months prior to

[14]Annualized housing expenditures equal the community price times the number of homogenous units of housing services consumed. Or,

$$r_{jtn} = p_j \times q(z_{jtn}),$$

where $q(z_{jtn})$ denotes the quantity index that translates heterogenous characteristics of housing unit n located in community j and rented in year t into equivalent homogenous units of housing services. This relationship is robust to a log transformation,

$$\ln(r_{jtn}) = \ln(p_j) + \ln(q(z_{jtn})).$$

If we define $\ln((q((z_{jtn}))) = \beta' S_{jtn} + \sum_t \delta_t Y_{tn} + \epsilon_{jtn}$, then the jth community fixed effect, δ_j's represents the log price index for community j.

[15]The EPA is required to announce changes in the status of the Superfund site 60 days prior to the official change in status; this information is published in the Federal Register.

the official cleanup date. Tables C.6 contain the results from these runs. A comparison of parameter estimates suggest that anticipation has no effect on the community price.

4.4.3 Superfund Sites

Hazardous waste sites are a disamenity for local housing. Toxic contaminants, possible odor, and visual blights make living near Superfund sites undesirable. A natural measure for this disamenity is a distance metric. Distance and inverse distance of housing units from Superfund sites have been used in several valuation studies to identify the price effect of proximity (Gamper-Rabindran et al. (2011), Kiel and Williams (2006), Kiel and Zabel (2001) among others). Kiel and Zabel (2001) find that price effects fall off sharply after three miles, which make community damages very localized.[16] While a direct measure of damage may seem more sensible than distance, an accurate measure is difficult to acquire. Damages vary by site and depend on unpredictable variation in weather patterns like wind and rainfall. The HRS score appears to be a likely candidate for a "damage index," however, it is a measure of the relative risk assessment by the EPA for determining the priority of cleanup, rather than a measure of actual damages (U.S. Environmental Protection Agency (2011b)).

The Euclidean distance of each house in the housing markets database to the seven Superfund sites is calculated using ArcGIS. Additional characteristics for these sites is acquired from the EPA. This data includes HRS scores, pollution pathways, etc. Figure C.2 displays the location of Superfund sites across Maricopa County.

To capture the nonlinear effects of proximity, we focus on the inverse distance. Rather than include all seven sites, we assume that households only care about proximity to the nearest site that has not completed cleanup. Each school district is assigned the transaction-weighted average inverse distance to the nearest Superfund site. Then cleanup or discovery of a site corresponds to a (possible) change in the inverse distance to the nearest site.

[16]Community damages can be large or small depending on the size of the community.

4.4.4 School Quality, Air Quality, and Urban Amenities

To control for other community-specific amenities, data on school and air quality is included in the public good index. Table C.7 summarizes the measures for school and air quality. Output-based measures of school quality are used in the analysis.[17] The average scaled scores for the Math, Reading, and Writing sections of Arizona's Instrument to Measure Standards (AIMS) exam is collected from School District Report Cards (Arizona Department of Education (2011)). Test scores are scaled between 500 and 900 points. Based on these adjusted scores, predefined cut points determine whether a student falls far below, approaches, meets, or exceeds the Arizona Department of Educations' standard. In addition to test scores, District Report Cards provide the percent of students in each school district that meet and/or exceed the state standard.

Air quality varies across Maricopa County. Data on PM_{10} is collected from the EPA for each Census tract in the county.[18] PM_{10} refers to particles less than 10 micrometers in diameter and includes a mixture of chemicals, metals, dust, among other things. According to U.S. Environmental Protection Agency (2011a), particulate pollution can lead to a variety of health issues dealing with respiration, such as irritation of the throat, coughing, and aggravated asthma. In some case, exposure may lead to death in people with lung or heart disease.[19]

In the public good index, school quality is measured as the percent of students that exceed the AIMS standard (averaged across the three sections of the exam). Due to restrictions on the availability of school quality data, the same data is used in the estimation of Models I and II. Air quality is measured as the transaction weighted annual mean PM_{10} level in mg/m^3 (a rescaling of the data for the estimation). An urban dummy variable is also included to control for access to local urban amenities, e.g. hospitals, parks, shopping centers, etc. The population

[17]Output measures of school quality are generally preferred to input measures such as student-teacher ratios and classroom dollars. In our dataset, the correlation between the price ranking of school districts and the student-teacher ratio is only 0.09, while the correlation between this price ranking and the average percent of students that exceed the AIMS standard is 0.85.

[18]The EPA interpolates air quality data collected from monitoring stations across Maricopa County. The distance from monitoring stations to the centroid of Census tracts vary significantly across the county, ranging from 3 miles to over 50 miles.

[19]The National Ambient Air Quality Standards (NAAQS) sets a maximum level for various air pollutants. For PM_{10}, the average level over a 24-hour period cannot exceed 150 $\mu g/m^3$ (U.S. Environmental Protection Agency (2011a)). In Maricopa County, the number of exceedences of this standard varies greatly. See Table C.7 for description of the variability in air quality across communities.

density varies significantly across Maricopa County, from 1.9 people per square mile to over 7000. Following the definition used by the U.S. Census Bureau for the assignment of Census tracts as urban/rural, the dummy equals 1 if the population density of the school district exceeds 1000 people per square mile.

4.4.5 Household Characteristics

Models I, II, and III require information on the characteristics of households (income and population) across Maricopa County. The 2000 U.S. Decennial Census provides publicly available spatial information on these household characteristics.[20] As the 2000 Census represents demographic information collected in 1999, it is appropriate for use in Model I. The Census Bureau, however, only collects this data once every 10 years. To fill in missing information for Models II and III, additional data was purchased from Geolytics. Geolytics is a private firm that compiles and maintains spatial demographics data for sale to marketing firms and researchers. A custom run of annual estimates on population and income at the 2000 Census block group level was performed.

The Census Bureau classifies income into 16 bins. The lowest income bin contains a count of households with an annual income less than $10,000. The top bin is unbounded with a count of households with annual income that exceeds $200,000. In the following, midpoint values for bins are assigned to the corresponding households to generate a distribution on income. The upper income bin is bounded at $500,000 (approximately 10 times the median household income in Maricopa County).[21]

4.4.6 Interpolation

The majority of data collected for estimation of the sorting model is at the Census tract level. Before estimation can begin this data must be interpolated into school districts.[22] Spatial weights were collected from the Missouri Census Data Center. This center provides spatial

[20]The data used in the estimation comes from the 100% sample SF3 file.

[21]Alternative methods exist for inferring a distribution from bin data, such as the censored interval regression. An exploration of these methods is left for future work.

[22]Data interpolation introduces additional measurement error in the estimation of the sorting model.

weights based on land area, density of population, and the density of housing units across a defined landscape. Additional weights were constructed using ArcGIS.

Different weights are required for count and continuous data. To illustrate, consider data collected from Census tract A that is be interpolated onto school district B; as depicted in Figure C.5. If data is continuous (such as a measure for air quality), then the appropriate weight is the proportion of the school district that contains the Census tract ($c/(b+c)$ in the figure). Weights calculated in this fashion will sum to 1 and create a weighted average for the continuous variable. Alternatively, if the data is a count variable (such as population), then the appropriate weight is the proportion of the Census tract that contains the school district ($c/(a+c)$ in the figure). These weights represent the share of the population in A that is contained in B and therefore need not sum to 1.

4.4.7 Joint Estimation

Joint estimation relies on the same structure of the sorting model presented in Section 4.3, and recovers parameter estimates through simulated GMM. In this case, the GMM moments stack information from both before and after the discrete change, which allows unobservables to change over time. By restricting parameters to be time-invariant, Model III estimates offer a robustness check for the parameter estimates of Models I and II. To maintain the normalization of prices across states of the world, community prices used in the joint estimation are relative to the price in the lowest community in the ex ante state. Differences in relative prices across states, not explained by observed changes in the public good, are attributed to variation in unobservable components over time.

4.5 Results

In the following, parameter estimates, inside and outside sample tests, and welfare analysis are examined for Models I, II, and III.

4.5.1 Estimation Results

Table C.8 contains parameter estimates for Models I - III.[23,24] In all three models, the parameter estimates are statistically significant and have the expected signs. In addition, the parameter values are similar to previous studies in the sorting literature. A few features deserve comment.

First, in all three models $\rho < 0$, which supports the underlying theory of the sorting model (that the single-crossing condition is satisfied). Second, the results suggest that households are more price elastic in Model II than Model I. This is in line with the housing boom that Maricopa County (and much of the U.S.) experienced in the first half of the 2000's. As noted by Pryce (2001), households are more responsive to changes in price during a housing boom, since the number of housing units available on the market increases and therefore so does the likelihood of finding lower priced alternatives. Third, the results imply that households place relatively less importance on school quality in Model II. The coefficients on the public good index represent the marginal rate of substitution with school quality. In Model II, the absolute value of parameter estimates for air quality and the inverse distance to the nearest NPL site are larger than in Model I. This suggests a shift of preferences away from education to these environmental amenities. One caveat to this interpretation is that data restrictions forced estimation of Models I and II using the same school quality data.[25] If school quality improved between Models I and II, then this would overstate the relative importance of the other public good components.

Table C.9 displays the impact of changes in community-specific attributes on the public good index. To control for differences in the scale and units of these attributes, one standard deviation changes are used. As is clear from the table, given the parameter estimates of the models, school quality dominates the provision of the public good. As such, even with the cleanup of a Superfund site, the effect on the public good index is relatively small. One possible

[23] Preliminary analysis found multiple local equilibria. A global search, using a genetic algorithm (run for one week), was performed to determine "good" starting values.

[24] Standard errors are calculated using simulation. See Berry et al. (2004) for a detailed discussion of parameter estimation in differentiated product demand systems.

[25] This implicitly assumes that school quality is exogenous to the sorting and constant across time. See Klaiber and Smith (2010) for a treatment of public education quality as endogenous to the sorting of households.

explanation for this is that households do not care about hazardous waste sites. Alternatively, this could be evidence of a poor choice of community definition. If the disamenity is very localized, as found in Kiel and Zabel (2001), then using school districts may "wash out" the impact on households living near the site.[26] Therefore, the cleanup and discovery scenarios that follow can be thought of as a small change in the public good.

4.5.2 Inside Sample Validation

Estimation of the model parameters imposes structure on the distribution of log income and log housing expenditures. The purpose of inside sample validation is to see how well the data fits this structure. In the following, we focus on the deviation and absolute deviation between fitted and empirical distributions as a measure of in-sample fit. We analyze this "fit" at quantiles along the entire distribution and for specific school districts. Table C.10 summarizes these results, which are averaged across the J communities and separated by quantiles.[27]

As Table C.10 indicates, averaged across communities and quantiles, the fitted distribution overstates log income by an average of 0.02% (0.02%) and understates log housing expenditures by an average of −0.07% (−0.39%) in Model I (II). However, this measure of fit can be misleading. Large positive and negative errors can cancel each other out and create a small average error. Using the absolute deviation between the fitted and empirical distributions, we find that the fitted distribution generates a 2.17% (2.18%) bias on log income and a 1.82% (1.88%) bias on log housing expenditures in Model I (II). The log transformation understates the average size of the bias. Using the unadjusted values for income and housing expenditures, the bias is an order of magnitude larger; 24.47% (24.76%) average absolute deviation between fitted and observed income and 17.99% (18.19%) on housing expenditures in Model I (II).

There are a few reasons to expect a better fit on the distribution of housing expenditures. First, the spatial resolution of the data is superior. Each transaction in the dataset is matched

[26]For example, only 0.03% of housing transactions and 11% of the land area in Dysart Unified District occurred within a 3-mile ring of Luke Air Force Base. As such, the impact of cleanup, at the school district level, may be quite small.

[27]To reduce the error associated with using a finite empirical sample, communities with fewer than 100 housing transactions were dropped from the calculations in the table. For example, the community Gila Bend Unified School District has 19 (33) observations in Model I (II) producing a coarse empirical distribution on log housing expenditures.

to its latitude/longitude, providing an improved connection with locally provided public good attributes (such as distance to nearest NPL site). Second, there is measurement error in the income distribution from its approximation using censored interval data. Finally, as noted from a comparison of equations (4.9) and (4.10), there are more parameters in the moment condition for log expenditures than log income.

The average percent absolute deviation is a useful statistic for the overall in-sample fit of the data. The bias, however, is not uniform across these log distributions. As can be seen from the table, the average fit is worse in the tails and largest in the lower tail. In fact, the bias on log income at the 5th quantile is over three times as large as at the 95th quantile.

The in-sample fit varies across school districts. Consider Tolleson, Phoenix, and Tempe Union High School Districts, which are ranked 7, 12, and 21 in Model I (7, 15, and 20 in Model II). Figures C.6 - C.9 display the fitted and empirical distributions on log income and log housing expenditures for these communities. As is apparent from the figure, while the overall fit for log income in Tempe Union High School District is "close," the fit for Phoenix Union High School District is relatively "poor." This is a product of the method of estimation. The simulated GMM minimizes the average distance between these curves at the 25th, 50th, and 75th quantiles. That is, it fits to the log distributions of the average community rather than individual communities. In addition, the tails of the distribution are ignored, which may be important depending on the skew of the log normal density. Finally, the model has only a few parameters that can be adjusted to improve the fit between the estimated and empirical distributions. This can be especially difficult as the empirical distributions are not constrained to the properties of the log normal density and parameter adjustment must come at the expense of fitting the data along another dimension.

Epple and Sieg (1999) also perform an in-sample test, but find a more favorable fit with their data. This result is partly explained by their method used to recover model parameters and the design of their in-sample test. Epple and Sieg employ a two-stage estimator that uses simulated GMM to match log income quantiles (25th, 50th, and 75th) in the first stage and nonlinear least squares on a boundary indifference equation (similar to equation (4.8)) in the second stage. Two assumptions in their analysis, when taken together, suggest that the in-sample test is misleading. First, communities are rank ordered by median income with price indices estimated to match this rank ordering. One concern with this strategy is that the ascending bundles property implies that conditional on taste, median income (or any other quantile) and price move together. This does not necessarily hold for unconditioned income because of the negative correlation between taste and income. Figures C.10 and C.11 illustrate the potential difference in ranking communities by median income and community prices, using our dataset. If the ascending bundles property holds for the unconditional mean, it should be the case that the points (triangles in the figure) should lie along a 45 degree line; the dashed line in the figure. As can be seen, none of these points do. Second, the authors do not look at the fit along the distribution of income. Instead, they focus on the fit for communities at the quantiles used in the estimation (Figure 6 in their paper). Not surprising, the predicted and estimated income move together in this setting.

4.5.3 Outside Sample Validation

The pure-characteristics vertical sorting model imposes substantial structure on the substitution patterns of households in response to a change in the provision of the public good. This structure defines how households react to changes in the public good, which housholds make new location choices, and which school districts they choose for these new location choices. The purpose of outside sample validation is to test this sorting mechanism. That is, how well does the sorting mechanism capture the observed substitution patterns of households. Two prospective analysis scenarios are considered, and provide a measure of *out-of-sample* fit. First, the parameter estimates of Model I are used to predict the re-sorting of households in response to the cleanup of Luke Air Force Base. Second, the mirror case is examined using the parameter estimates of Model II to predict the response to the discovery of hazardous waste at this same military base. Given the changes proposed in each scenario, the corresponding counterfactual equilibrium is computed (as outlined in Section 4.3.6).[28] The computed data on income, housing expenditures, community prices, and population shares is compared with the data observed

[28] Along with the proposed change, observed changes in air quality are included in the out-of-sample comparison.

from the quasi-experiment.

Table C.11 summarizes the results for the out-of-sample fit on the log distributions for income and housing expenditures. Again, these results are averaged across the J communities and broken out by quantiles. Comparing the absolute deviation between the fitted and empirical distributions, we find a 2.15% (2.12%) bias on log income and a 1.93% (2.01%) bias on log housing expenditures in Model I (II). Using the unadjusted values, the bias is an order of magnitude larger; 24.27% (24.03%) bias on income and 17.77% (21.15%) on housing expenditures in Model I (II). These results are very similar to those in the in-sample fit, which is not surprising given the size of changes in the public good. Again, the bias is not uniform across the distribution of income/housing expenditures and is largest at the tails of the distribution. Figures C.12 through C.15 display the fitted and empirical distributions on log income and log housing expenditures for the same three school districts as the in-sample fit. Surprisingly, compared with the in-sample results for these school districts, the overall fit appears to improve.

In addition to the above bias, Model I (II) tends to under-predict (over-predict) community prices, as displayed in Figures C.16 and C.17. The vertical sorting model is not designed to capture speculation in the housing market.[29] It is a static model with households that rent housing units from absentee landlords. As such, the sorting model overlooks the housing boom between Models I and II. Market speculation (during the housing boom) plus changes in unobservable community attributes change the rank ordering of communities, which is missed in the prediction of Models I and II; see Table C.5.[30] Speculation is highlighted in Figure C.20, which displays the price/public good relationship (from Model III estimates) pre- and post-cleanup of Luke Air Force Base. Though the public good does not change much between these periods, prices rise, a vertical shift of the curve in the figure, consistent with forward-looking speculative behavior.[31]

One concern is that the vertical sorting model overstates the response of households because moving is costless. In both scenarios the computed sorting of households appear to provide

[29]Bayer et al. (2010) incorporate dynamic behavior into sorting.

[30]Neither Model I or II predict any change in the rank ordering of school districts.

[31]See Morris (1996) and Harrison and Kreps (1978) and for a discussion of speculative premiums in asset markets.

a good approximation for the observed population shares. Across the 24 communities, the maximum absolute deviation was only 0.004 (0.004) in Model I (II). In percentage terms, it is considerably higher 11.9% (13.6%) as some school districts hold a very small share of the population of Maricopa County, e.g. Agua Fria Union High School District contains 0.04% of the population. Figures C.18 and C.19 display the percent deviation for the 24 communities.

4.5.4 Welfare

Welfare analysis provides additional evidence for model validation. Though we do not observe the true welfare effect of a prospective policy, economic intuition provides guidance whether the estimated welfare effects are "reasonable." In the following, we focus on the two scenarios from the outside sample validation and examine the indirect implications of the sorting mechanism on welfare measure. That is, we begin with a standard welfare analysis for the scenarios and then examine the price and welfare effects on the direct substitutes (by ranking) of Dysart Unified District, the school district containing Luke Air Force Base. Tables C.12 and C.13 contain the annualized mean willingness-to-pay results separated by school district and household type (movers, nonmovers, etc).

Let us begin with the case of cleanup of Luke Air Force Base (using Model I parameter estimates). The mean partial equilibrium willingness-to-pay (WTP_{PE}) is small, but positive for the school districts that experience an improvement in the public good. WTP_{PE} values range from \$0 to \$0.17 across school districts. Figure C.21 displays the spatial distribution of WTP_{PE}. Note that the highest welfare values are for the sites that directly experience a change from the site cleanup (those with Luke Air Force Base as the nearest Superfund site). School districts further away, or which have other Superfund sites closer, do not experience a welfare improvement. For example, Luke Air Force Base is near the border of Agua Fria Union High School District. In fact, a 3-mile ring surrounding the Superfund site contains some of this school district. Because of the assumption that only the nearest site matters to households, cleanup of Luke Air Force Base has no estimated welfare effect on this school district. Using population shares to weight the mean values generates a mean WTP_{PE} of \$0.01 for Maricopa County, an annualized WTP_{PE} of approximately \$30,000.

The general equilibrium welfare measure incorporates the response of households and the market to changes in the public good. In this setting, households can experience a welfare loss from an improvement in the public good because of changes in relative prices. A description of the general welfare results is complicated by the fact that different households in a school district experience different changes in prices and public goods because of movement (Sieg et al. (2004)). For comparability with the partial equilibrium measure, consider the mean general equilibrium willingness-to-pay (WTP$_{GE}$) of households that originated in a school district. Compared with the partial equilibrium results, mean WTP$_{GE}$ spans a much wider range of values; the maximum mean welfare loss is $-\$52.78$ (Aguila Elementary District) and maximum welfare gain is \$24.50 (Nadaburg Elementary District). The population weighted mean willingness-to-pay is still quite small, \$0.19, with an annualized WTP$_{GE}$ of approximately \$580,000.

Averaged across Maricopa County, mean partial and general equilibrium WTP are similar.[32] The potential strength of the sorting model is not in mean welfare analysis. Rather, it is the ability of the model to study the distributional consequences of government policy. That is, who are the "winners and losers" to a policy change. Figure C.22 displays the spatial distribution of mean WTP$_{GE}$ for the households originating in these school districts. Unlike the partial equilibrium measure, the gains and losses are more pronounced and spread across the entire county. For example, in the partial equilibrium setting Tempe Union High School District has a mean WTP$_{PE}$ of \$0. Since other Superfund sites are closer to this school district and household location choices and prices are fixed, the cleanup has no partial equilibrium welfare effect. In a general equilibrium setting, cleanup induces households to re-sort. As households move into Tempe, this places pressure on the housing market, which adjusts with higher prices. In the new sorting equilibrium, the households that lived in Tempe Union High School District prior to cleanup experience a mean welfare loss of $-\$1.53$ due to these price effects.

The sorting mechanism imposes structure on the spatial distribution of welfare estimates, as it defines the substitutes for households living in a school district. To illustrate, consider Dysart Unified District, which is ranked 15 and experiences the largest increase in the public good from cleanup. According to Table C.5, Higley Elementary School District (rank 14) and

[32] Statistical tests can evaluate whether the differences are significant.

Gilbert Unified District (rank 16) are the closest substitutes for the households living in Dysart Unified District. Given the model structure, cleanup of Luke Air Force Base attracts the wealthiest households (conditioned on taste) and those with the strongest taste for the public good (conditioned on income) from Higley Elementary School District into Dysart Unified District. The influx of households leads housing markets to adjust prices upward. But this changes the relative prices between Dysart Unified District and Gilbert Unified District, which induces the wealthiest households (conditioned on taste) and those with the strongest taste for the public good (conditioned on income) to move from Dysart Unified District into Gilbert Unified District. Again housing markets adjust and create new prices. This adjustment process continues with households moving back and forth until the model reaches a sorting equilibrium.

In terms of welfare, this means that households originally living in Dysart Unified District experience different changes in prices and the public good. As denoted in Table C.12, the households that move into Gilbert experience an increase in prices but also an increase in the public good provision. The net effect is a mean welfare improvement (mean WTP_{GE}=\$0.66). The households that remain in Dysart experience a fall in prices, as more households exit the school district than enter, and an increase in the public good. This is an clear welfare improvement (mean WTP_{GE}=\$0.73).

The welfare results for the discovery of hazardous waste at Luke Air Force Base are similar. Using Model II parameter estimates, mean WTP_{PE} range from $-\$0.47$ to \$0, while mean WTP_{GE} range from $-\$22.82$ to \$40.68. The population weighted mean WTP_{PE} and mean WTP_{GE} are $-\$0.02$ and $-\$0.26$, respectively. Differences between the results and those presented above are due to differences in parameter estimates. Figure C.23 displays the spatial pattern of mean WTP_{GE} for the households originally in a school district. Note that the spatial pattern for Model II is different from Model I (Figure C.22). This is due to the differences in rank ordering implied by the estimated price indices. Because of the different rank ordering, there is a different substitution pattern.

A comparison of the results also highlights the sensitivity of household mobility to the parameter estimates. As illustrated in Tables C.12 and C.13, Model II predicts more movement by households than Model I; a consequence of the parameter estimates for the price elasticity of housing demand and the CES substitution parameter. These parameters enter the cuts points (in equation (4.6)) that define the movement of households. In Model II, households are more sensitive to changes in housing prices (the absolute price elasticity is larger). In addition, the CES substitution parameter is also smaller in Model II. As such, the private and public good are weaker substitutes. Combined, these two components increase the movement of households in Model II.

4.6 Conclusions

In this paper we perform inside sample and outside sample validation exercises on the pure-characteristics vertical sorting model. Using the cleanup of Luke Air Force Base as a quasi-experiment, data surrounding the cleanup is used to compare the observed and predicted sorting of households, prices, income, and housing expenditures. Welfare analysis for cleanup and discovery of contaminants at Luke Air Force Base is included to highlight a proposed strength of the sorting model and provide additional evidence for validation. In-sample tests compare the fitted and empirical distributions on income and housing expenditures to examine how well the data fits the structure of the model.

Based on the results from the validation exercises, the predictive capability of the vertical sorting model does not come without a cost. We find an average bias (absolute deviation) of approximately 20% on housing expenditures and nearly 25% on income. This bias is not uniform across the distribution, and increases in the tails. The "fit" on individual school districts also varies widely. This finding is consistent across in-sample and out-of-sample tests. Though the model appears to fit population shares, it under-predicts prices in Model I and over-predicts in Model II. The movement of households is also sensitive to parameter estimates. The spatial distribution of welfare is a feature of the vertical sorting model's structure as the substitutes (school districts) are defined by the model. Mean general equilibrium welfare estimates vary by school district, ranging from $-\$52.78$ to \$24.50 for cleanup at Luke Air Force Base. Averaged across the school district populations, the mean general equilibrium willingness-to-pay was slightly higher than the partial equilibrium counterpart, \$0.19 and \$0.01 respectively.

These results come with a note of cation. Though cleanup of Luke Air Force Based offered a convenient quasi-experiment, it also occurred during a housing boom, which the model is not designed to capture. Also, there is evidence that the definition of community used in the paper may not be optimal for this setting, as the spatial effects of Superfund sites are localized. All this calls for additional efforts for model validation.

Further research will allow more definitive statements about the pure-characteristics vertical sorting model. Model validation is a Bayesian process that requires updating and repetition (Keane (2010)). An exploration of alternative definitions for the community, and its effects on the inside and outside sample fit of the model, would generate new insights. One possible alternative is to construct artificial communities, similar to Banzhaf and Walsh (2008), to capture a higher spatial resolution of household movement. However, such assumptions come at the cost of an economic rationale.

In addition, data from the quasi-experiment contains information about the movement of households and the response of the housing market since the initial inclusion of sites on the NPL list. A useful model validation exercise would compare the inside and outside sample fit of the model from initial discovery through present day remediation. As such, Maricopa County would provide multiple opportunities to examine the sorting model.

Bibliography

Arizona Department of Education, 2011. District Report Cards. Available online at: http://http://www10.ade.az.gov/reportcard/.

Banzhaf, H. S., Walsh, R. P., June 2008. Do People Vote with Their Feet? An Empirical Test of Tiebout's Mechanism. American Economic Review 98 (3), 843–863.

Bayer, P., Keohane, N., Timmins, C., 2009. Migration and Hedonic Valuation: The Case of Air Quality. Journal of Environmental Economics and Management 58 (1), 1–14.

Bayer, P., McMillan, R., Murphy, A., Timmins, C., 2010. A Dynamic Model of Demand for Houses and Neighborhoods. Working Paper.

Bayer, P., McMillan, R., Reuben, K., 2004. An Equilibrium Model of Sorting in an Urban Housing Market. NBER Working Paper No 10865.

Berry, S., Linton, O., Pakes, A., 2004. Limit Theorems for Estimating the Parameters of Differentiated Product Demand System. The Review of Economic Studies 71 (3), 613–654.

Berry, S., Pakes, A., 2007. The Pure Characteristics Demand Model. International Economic Review 48 (4), 1193–1225.

Ellickson, B., 1971. Jurisdictional Fragmentation and Residential Choice. American Economic Review 61 (2), 334–339.

Epple, D., Filimon, R., Romer, T., 1993. Existence of Voting and Housing Equilibrium in a System of Communities without Taxes. Regional Science and Urban Economics 23, 585–610.

Epple, D., Filimon, R., Romer, T., 1994. Equilibrium Among Local Jursidictions: Towards an Integrated Approach of Voting and Residential Choice. Journal of Public Economics 24, 281–304.

Epple, D., Platt, G. J., 1998. Equilibrium and Local Redistribution in an Urban Economy when Households Differ in both Preferences and Income. Journal of Urban Economics 41 (1), 23–51.

Epple, D., Romer, T., 1991. Mobility and Redistribution. Journal of Political Economy 99 (4), 828–858.

Epple, D., Romer, T., Sieg, H., 2001. Interjurisdictional Sorting and Majority Rule an Empirical Analysis. Econometrica 69 (6), 1437–1465.

Epple, D., Sieg, H., 1999. Estimating Equilibrium Models of Local Jurisdiction. Journal of Political Economy 107 (4), 645–681.

Freeman, A. M., 2003. The Measurement of Environmental and Resource Values. RFF Press.

Gamper-Rabindran, S., Mastromonaco, R., Timmins, C., 2011. Valuing the Benefits of Superfund Site Remediation: Three Approaches to Measuring Localized Externalities. NBER Working Paper Series, Working Paper 16655.

Harrison, M., Kreps, D., 1978. Speculative investor behavior in a stock market with heterogeneous expectations. The Quarterly Journal of Economics 92 (2), 323–336.

Keane, M., 2010. Structural vs. Atheoretic Approaches to Econometrics. Journal of Econometrics 156 (1), 3–20.

Kiel, K. A., Williams, M., 2006. The Impact of Superfund Sites on Local Property Values: Are All Sites the Same? Journal of Urban Economics (61), 170–192.

Kiel, K. A., Zabel, J., 2001. Estimation the Economics Benefits of Cleaning Up Superfund Sites: The Case of Woburn. Journal of Real Estate Finance and Economics (22), 163–184.

Klaiber, H., Phaneuf, D., 2010. Valuing Open Space in a Residential Sorting Model of the Twin Cities. Journal of Environmental Economics and Management 60 (2), 57–77.

Klaiber, H. A., Smith, V. K., 2010. General Equilibrium Benefit Analyses for Social Programs. Working Paper.

Kuminoff, N. V., 2009. Decomposing the Structural Identification of Nonmarket Values. Journal of Environmental Economics and Management 57 (2), 123–139.

Kuminoff, N. V., Smith, V. K., Timmins, C., 2010. The New Economics of Equilibrium Sorting and its Transformational Role for Policy Evaluation. National Bureau of Economic Research Working Paper # 16349.

Morris, S., 1996. Speculative investor behavior and learning. The Quarterly Journal of Economics, 1111–1133.

Porterba, J. M., 1992. Taxation and housing: old questions, new answers. Empirical Public Finance 82 (2), 237–242.

Pryce, G., 2001. Cycles in The Price Elasticity of Demand for Housing. University of Glasgow, Ecofin Discussion Paper No.5.

Sieg, H., Smith, V., Banzhaf, H., Walsh, R., 2004. Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods. International Economic Review 45 (4), 1047–1077.

Sieg, H., Smith, V. K., Banzhaf, S., Walsh, R., 2002. Interjurisdictional Housing Prices in Locational Equilibrium. Journal of Urban Economics 52, 131–153.

Smith, V., Banzhaf, H., 2004. A Diagrammatic Exposition of Weak Complementarity and the Willig Condition. American Journal of Agricultural Economics 86 (2), 455–466.

Smith, V., Banzhaf, H., 2007. Quality Adjusted Price Indexes and the Willig Condition. Economic Letters 94, 43–48.

Smith, V., Sieg, H., Banzhaf, H., Walsh, R. P., 2004. General Equilibrium Benefits for Environmental Improvements: Projected Ozone Reductions Under EPAs Prospective Analysis for the Los Angeles Air Basin. Journal of Environmental Economics and Management 47, 559–584.

Tiebout, C. M., 1956. A Pure Theory of Local Expenditures. Jourbnal of Political Economy 64 (5), 416–424.

Tra, C. I., 2010. A Discrete Choice Equilibrium Approach to Valuing Large Environmental Changes. Journal of Public Economics 94, 183–196.

U.S. Environmental Protection Agency, 2011a. Air and Radiation. Available online at: http://http://www.epa.gov/air/.

U.S. Environmental Protection Agency, 2011b. Superfund Program. Available online at: http://www.epa.gov/superfund/.

Walsh, R., 2007. Endogenous Open Space Amenities in a Locational Equilibrium. Journal of Urban Economics 61 (2), 319–344.

Westoff, F., 1977. Existence of Equilibria in Economies with a Local Public Good. Journal of Economic Theory 14 (1), 84–112.

Willig, R. D., 1978. Incremented Consumer's Surplus and Hedonic Price Adjustments. Journal of Economic Theory 17, 227–253.

CHAPTER 5. GENERAL CONCLUSIONS

5.1 Conclusions

The three essays in this dissertation focus on problems concerning uncertainty, learning, and welfare measurement in resource and environmental economics. The first two essays focus on the commercial fishery and extend our understanding of the behavior and management of these fishermen under a dynamic setting of choice with uncertainty and learning. The third essay provides the first external model validation for the pure-characteristics vertical sorting model, which is a general equilibrium structural econometric model capable of estimating spatial and distributional welfare effects from large-scale policy changes. The following describes the main conclusions from this body of work.

The first essay analyzes search, learning, and economic performance in a dynamic fishing game. The analysis contrasts equilibrium site choice strategies under a first-best benchmark, the case of independent, non-cooperative fishermen, and under a stylized fishing cooperative. This paper extends our understanding of the behavior of fishermen operating in a fishing cooperative and highlights the importance of the design of fishing policy to promote their use. Independent fishermen do not internalize the full value of information and do not replicate first-best search patterns. An information sharing cooperative, though an improvement relative to independent effort, faces a free-riding problem as each member prefers that the costly search for information be undertaken by others. Devising contracts that result in optimal investment in information may be particularly challenging in fisheries, due to the club good characteristic of information, its costly acquisition, and the common property nature of the fishery resource. The addition of risk aversion significantly alters the equilibrium pattern of search and the demand for information through the introduction of disutility associated with experiencing risk.

The second essay presents a dynamic transition model to study the economic fundamentals that determine the path from an initially over-capitalized fishery to the ITQ-regime fleet structure. The simulation results highlight the model implications regarding delayed-exit strategies, speculation, and the generation of rents. The results suggest that beliefs play a strong role on the transition. Uncertainty by fishermen over their relative cost-efficiency translates into uncertainty over quota trading prices. As such, a component of the ITQ asset's price is speculation over future trading prices. Heterogeneous uncertainty and learning also provide insight into observed patterns of vessel exit, generating new insight into the slow transition observed in U.S. fisheries. These results should be of interest to policy makers. The near universal practice of allocating the initial endowment of quota based on historic catch promotes delayed-exit strategies on cost-inefficient vessels and thereby prolongs the transition period during which the full efficiency benefits of ITQ management are unrealized.

The third essay uses the cleanup of Luke Air Force Base as a quasi-experiment to perform inside sample and outside sample validation exercises on the pure-characteristics vertical sorting model. Based on these validation exercises, the predictive capability of the vertical sorting model does not come without a cost. The average bias (absolute deviation) on the estimates for housing expenditures is approximately 20% and nearly 25% on income. This bias is not uniform across the distribution, and increases in the tails. The fit on individual school districts also varies widely. This finding is consistent across in-sample and out-of-sample tests. Though the model appears to fit population shares, it under-predicts prices in Model I and over-predicts in Model II. The movement of households is sensitive to parameter estimates. The spatial distribution of welfare is also a feature of the vertical sorting model's structure as the substitutes are defined by the model.

APPENDIX A. ADDITIONAL MATERIAL FOR CHAPTER 2

A.1 Belief Updating

In this section we demonstrate how fishermen use their private and shared (if available) signal(s) to update their beliefs about the payoffs at fishing sites; see Gelman et al. (2003) for deeper discussion of Bayesian statistics.

Independent fishermen observe a single payoff signal in a fishing period. Suppose the Row fisherman fishes site j in the first fishing period and observes the payoff signal s_j. Using Bayes' rule, his updated beliefs about site j are,

$$\mu_j'|s_j = \theta_j\mu_j + (1-\theta_j)s_j$$
$$\nu_j'|s_j = \theta_j\nu_j$$

where $\theta_j = \frac{\nu_s}{\nu_j+\nu_s}$ denotes the weight placed on prior beliefs. As the noise associated with the signal grows, more weight is placed on prior information and learning slows (as $\nu_s \to \infty$ we see $\theta_j \to 1$). The above is a simple reweighting of his prior belief with the new payoff information. As the fisherman does not observe any new information about site i, his beliefs remain unchanged. Using the provided updating rules, we can calculate $b' = \{\mu_1', \nu_1', \mu_2', \nu_2'\}$ for independent fishermen.

Fishermen within the cooperative receive two independent payoff signals in a fishing period. This creates two possible cases of interest, either the fishermen fished the same site or separate sites. Suppose the cooperative sends both fishermen to fish site j in the first fishing period and the Row fisherman receives payoff signals s_j and $\breve{s}_j$. Using Bayes' rule, his updated beliefs about site j are,

$$\mu_j' \,|s_j, \breve{s}_j \;=\; \theta_j \mu_j + (1-\theta_j)\bar{s}_j$$
$$\nu_j' \,|s_j, \breve{s}_j \;= \theta_j \nu_j$$

where $\theta_j = \frac{\nu_s}{2\nu_j + \nu_s}$ and $\bar{s}_j$ denotes the average payoff signal observed at site j. Again, this is a simple reweighting of his prior with the new payoff information. In contrast with the case of independent fishermen, the Row player places more weight on the new information, reflecting the use of two signals. As there is no new information about site i, his beliefs remain unchanged.

Now suppose the cooperative sends fishermen to separate sites in the first fishing period and the Row fisherman receives payoff signals s_j and $\breve{s}_i$. Using Bayes' rule, his updated beliefs are,

$$\mu_j' \,|s_j, \breve{s}_i \;= \theta_j \mu_j + (1-\theta_j) s_j$$
$$\nu_j' \,|s_j, \breve{s}_i \;= \theta_j \nu_j$$

where $\theta_j = \frac{\nu_j}{\nu_j + \nu_s}$, and for site i,

$$\mu_i' \,|s_j, \breve{s}_i \;= \theta_i \mu_j + (1-\theta_i) s_i$$
$$\nu_i' \,|s_j, \breve{s}_i \;= \theta_i \nu_i$$

Using the provided updating rules, we can calculate $b' = \{\mu_1', \nu_1', \mu_2', \nu_2'\}$ for the cooperative.

A.2 Bayesian Nash Equilibrium

In this section we solve the date $t = 1$ Bayesian game for independent fishermen in an expected utility framework. The case of risk neutrality is a special case ($\lambda \to 0$). A Bayesian Nash Equilibrium (BNE) requires players strategies be best responses to each other. For example, the Row player's strategy is a best response to the Column player if, for every possible signal observed by the Row player, the action specified by his action rule for that signal maximizes his expected payoff, given his belief about the Column players beliefs and the Column player's action rule.

Consider the following strategy profile, there may be others, for the Row and Column fisherman.

$$a'(s_j) = \begin{cases} \text{Fish site } j & \text{if } s_j > s^U \\ (a'_1,\ a'_2) & \text{if } s^L \leq s_j \leq s^U \\ \text{Fish site } i & \text{if } s_j < s^L \end{cases}$$

and

$$\breve{a}'(\breve{s}_j) = \begin{cases} \text{Fish site } j & \text{if } \breve{s}_j > \breve{s}^U \\ (\breve{a}'_1,\ \breve{a}'_2) & \text{if } \breve{s}^L \leq \breve{s}_j \leq \breve{s}^U \\ \text{Fish site } i & \text{if } \breve{s}_j < \breve{s}^L \end{cases}$$

We need to calculate the thresholds s^U, s^L, $\breve{s}^U$, and $\breve{s}^L$ that ensure the above contingent strategy profile satisfies a BNE. There are four possible scenarios to consider depending on where fishing occurred in the first period: both fishermen fished at site 1, both fished at site 2, the Row fisherman fished at site 1, while the Column fisherman fished at 2, and visa versa. We focus on the optimal policy and payoffs for the case where both fishermen fish site 1 in the first fishing period. The analysis of the remaining cases follow analogously and to conserve space are not repeated.

We begin by specifying how fishermen form beliefs about beliefs. At date $t = 1$, each fisherman has observed some private information and the site choice of the other. We assume that fishermen use their updated beliefs to form rational conjectures regarding the choices and beliefs of their counterpart. Since both fishermen fished site 1 in the first period, their belief about their rival's private information should be a function of their own private information.

$$b(\breve{s}_1) \sim N\left(\mu'_1,\ \nu'_1 + \nu_s\right)$$

$$\breve{b}(s_1) \sim N\left(\breve{\mu}'_1,\ \breve{\nu}'_1 + \breve{\nu}_s\right)$$

where $b(\breve{s}_1)$ denotes the Row player's belief about the Column player's signal.

Suppose $\breve{s}^U$ and $\breve{s}^L$ exist such that the Column player plays the above strategy; $\breve{a}'(\breve{s}_1)$. The Column player knows $\breve{s}_1$, but the Row player does not (he only has beliefs according to $b(\breve{s}_1)$). Then given the Column player's strategy, the Row player's expected utility from fishing site 1 or site 2 are,

$$
\begin{aligned}
EU[\text{site 1}] =& EU\Big[\Pr\left(\breve{s}_1 > \breve{s}^U|s_1\right)\left(s_1' - \kappa\right) + \Pr\left(\breve{s}_1 < \breve{s}^L|s_1\right)s_1' \\
&+ \Pr\left(\breve{s}^L \leq \breve{s}_1 \leq \breve{s}^U|s_1\right)\left[\breve{a}_1'(s_1' - \kappa) + (1 - \breve{a}_1')s_1'\right]\Big] \\
EU[\text{site 2}] =& EU\Big[\Pr\left(\breve{s}_1 > \breve{s}^U|s_1\right)s_2' + \Pr\left(\breve{s}_1 < \breve{s}^L|s_1\right)\left(s_2' - \kappa\right) \\
&+ \Pr\left(\breve{s}^L \leq \breve{s}_1 \leq \breve{s}^U|s_1\right)\left[\breve{a}_1's_2' + (1 - \breve{a}_1')\left(s_2' - \kappa\right)\right]\Big]
\end{aligned}
$$

For the $s_1 \in [s^L,\ s^U]$ type Row player to mix his two actions, $a_1' \in (0, 1)$, it must be the case that based on his conjecture about the Column player's strategy $\breve{a}'(\breve{s}_1)$, he is indifferent between fishing site 1 and site 2. Setting the expected utility of fishing site 1 and site 2 equal, conditional on the Column players contingent strategy, and solving for $\breve{a}_1'$, we find,

$$
\breve{a}_1'(s_1) = \frac{\mu_1' - \mu_2' - \frac{\lambda}{2}(\nu_1' - \nu_2') - \kappa + 2\kappa\Pr\left(\breve{s}_1 < \breve{s}^U|s_1\right)}{2\kappa\Pr\left(\breve{s}^L \leq \breve{s}_1 \leq \breve{s}^U|s_1\right)} \quad \forall\ s_1 \in [s^L,\ s^U] \tag{A.1}
$$

Using the same method, we can calculate the expected utility of fishing either site for the Column player. For the $\breve{s}_1 \in [\breve{s}^L,\ \breve{s}^U]$ type of the Column player to mix his two actions, $\breve{a}'(\breve{s}_1) \in (0, 1)$, it must be the case that based on his conjecture about the Row player's strategy $a'(s_1)$, he is indifferent between fishing site 1 and site 2. Again, we find,

$$
a_1'(\breve{s}_1) = \frac{\breve{\mu}_1' - \breve{\mu}_2' - \frac{\lambda}{2}(\breve{\nu}_1' - \breve{\nu}_2') - \kappa + 2\kappa\Pr\left(s_1 < s^U|\breve{s}_1\right)}{2\kappa\Pr\left(s^L \leq s_1 \leq s^U|\breve{s}_1\right)} \quad \forall\ \breve{s}_1 \in [\breve{s}^L,\ \breve{s}^U] \tag{A.2}
$$

In a BNE, each player's action must be optimal subject to their Bayesian updated belief about the strategy of their rival. Using equations (A.1) and (A.2) we derive four equations that define the parameters s^U, s^L, $\breve{s}^U$, and $\breve{s}^L$ that will satisfy a BNE.

The Row player believes the Column player is indifferent between mixing and fishing site 1 when $\breve{a}_1'(s_1)$ exactly equals one (the Column player knows that this is the Row player's belief). For site 1 to be strictly dominant, it must be that it is still preferred even under the worst case (when the Row player also wants to fish the site). Using this fact, the updating rules in A.1, and given the strategy profile, the Row player's upper signal threshold s^U and the Column player's upper signal threshold $\breve{s}^U$ must satisfy,

$$
s^U = \frac{1}{1 - \theta_1}\left\{\left[\mu_2 - \theta_1\mu_1\right] - \frac{\lambda}{2}\left[\nu_2 - \theta_1\nu_1\right] + \kappa\right\} \tag{A.3}
$$

The Row player believes the Column player is indifferent between mixing and fishing site 2 when $\breve{a}_1'(s_1)$ exactly equals zero (the Column player knows that this is the Row player's belief). For site 2 to be strictly dominant, it must be that it is still preferred even under the worst case (when the Row player also wants to fish the site). Using this fact and given the strategy profile, the Row player's lower signal threshold s^L and the Column player's lower signal threshold $\breve{s}^L$ must satisfy,

$$s^L = \frac{1}{1-\theta_1}\left\{\left[\mu_2 - \theta_1\mu_1\right] - \frac{\lambda}{2}\left[\nu_2 - \theta_1\nu_1\right] - \kappa\right\} \tag{A.4}$$

Using the same argument, we construct the following for the Column player,

$$\breve{s}^U = \frac{1}{1-\breve{\theta}_1}\left\{\left[\breve{\mu}_2 - \breve{\theta}_1\breve{\mu}_1\right] - \frac{\lambda}{2}\left[\breve{\nu}_2 - (1-\breve{\theta}_1)\breve{\nu}_1\right] + \kappa\right\} \tag{A.5}$$

$$\breve{s}^L = \frac{1}{1-\breve{\theta}_1}\left\{\left[\breve{\mu}_2 - \breve{\theta}_1\breve{\mu}_1\right] - \frac{\lambda}{2}\left[\breve{\nu}_2 - (1-\breve{\theta}_1)\breve{\nu}_1\right] - \kappa\right\} \tag{A.6}$$

The thresholds that satisfy the BNE , s^U, s^L, $\breve{s}^U$, and $\breve{s}^L$, defined by equations (A.3)-(A.6), must simultaneously hold.

A.3 Figures

Figure A.1 Chronology of Site Choice Problem

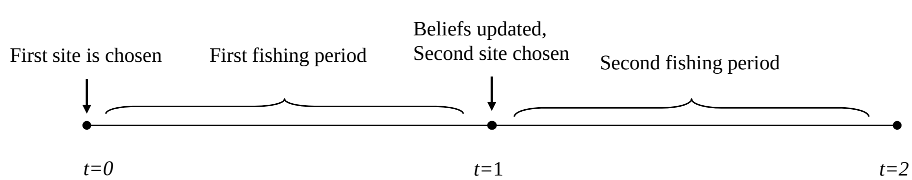

Figure A.2 Equilibrium Site Choices - Risk Neutrality

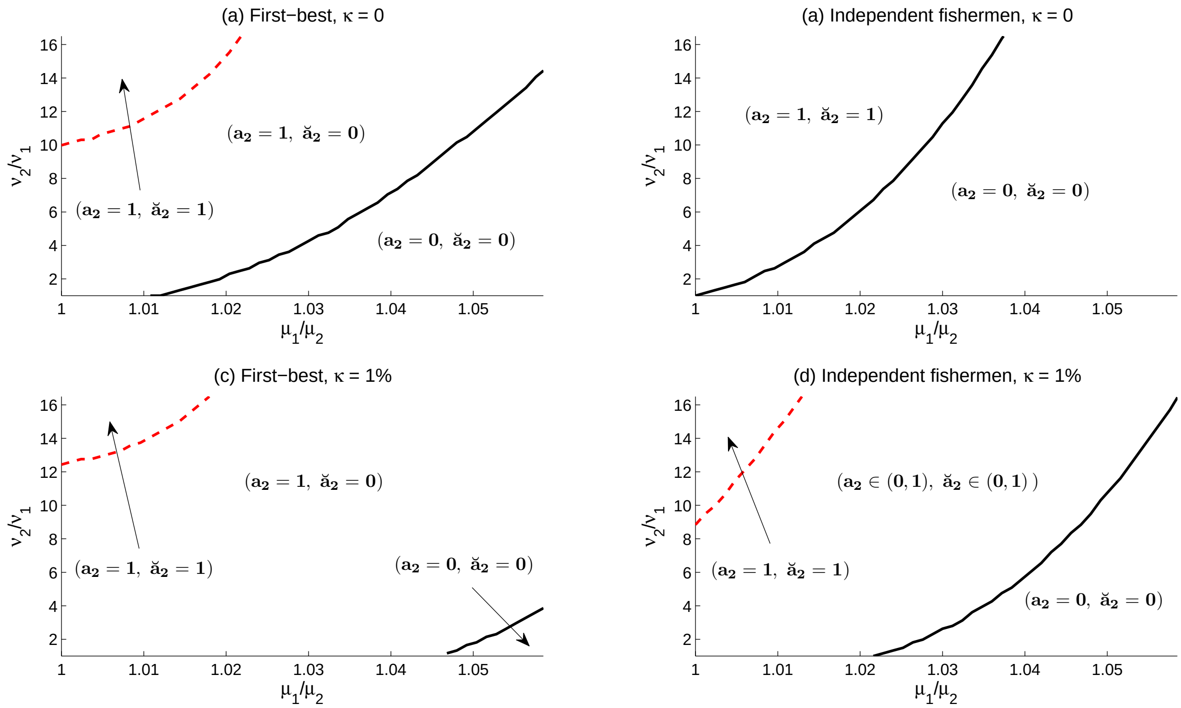

Figure A.3 Equilibrium Site Choices - Risk Aversion ($\lambda = 0.1$)

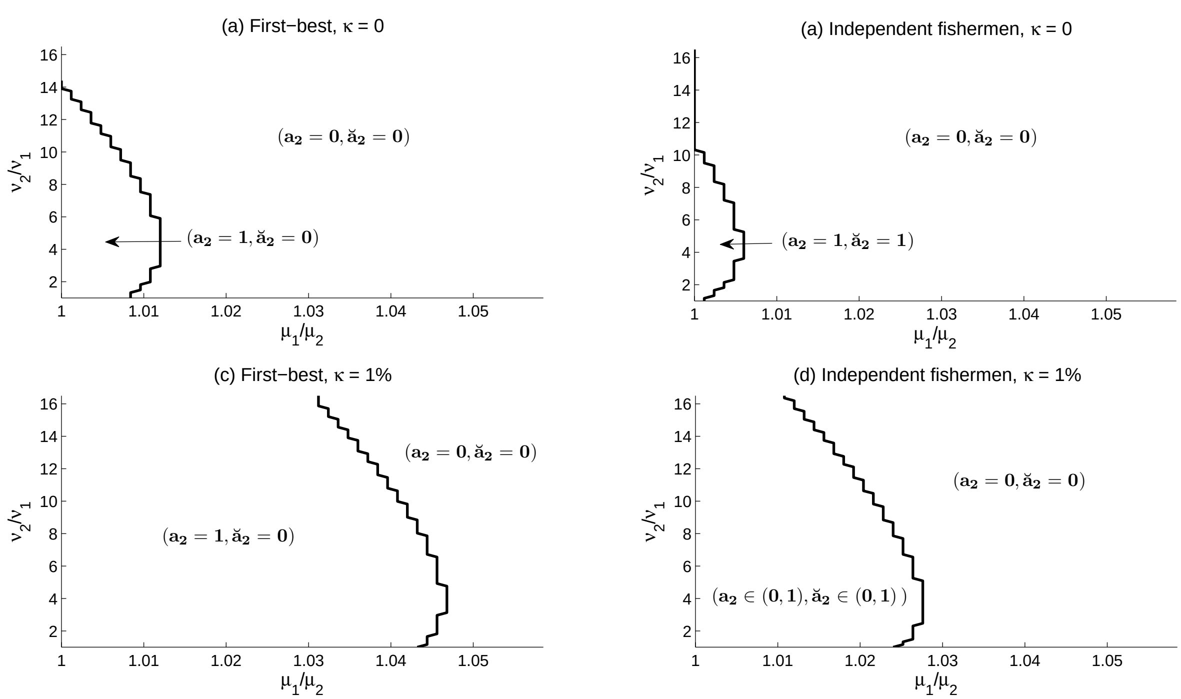

Figure A.4   Conditional Two-Period Expected Payoffs

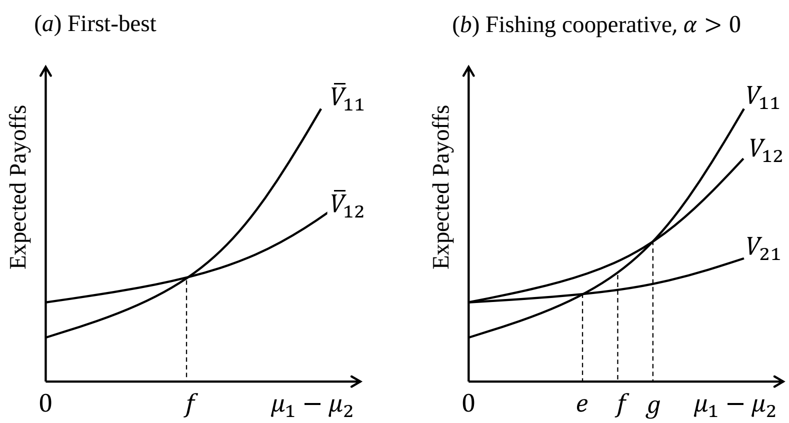

APPENDIX B. ADDITIONAL MATERIAL FOR CHAPTER 3

B.1 Learning

Each period $t = 1, \ldots, T$, the interaction of fishermen in the quota market generates noisy information about the distribution of skill among the population of fishermen. Fishermen use this public information to update their beliefs about the unknown mean skill. Ideally, the quota trading price would be the public signal, which would allow for the formation of asset bubbles within the model. However, for tractability, the public information in period t is modeled as an exogenous signal equal to the true mean plus some noise,

$$s_t = \beta + \epsilon$$

where ϵ is a mean zero normal random variable with small variance. This signal contains information similar to the quota trading price, but with a constant variance. An endogenous price signal would exhibit decreasing noise during the transition as some uncertainty is resolved each trading period.

Similar in form to the Bayesian learning rule for a normal prior, fishermen update their mean belief using the weighted average of their prior belief and the new information. Let μ_{t-1} and σ^2_{t-1} denote the mean and variance of beliefs contained in B_{t-1}, and θ the weight placed on prior beliefs. Then the updated mean belief in period t is,

$$\mu_t = \theta\mu_{t-1} + (1 - \theta)s_t.$$

Uncertainty is resolved in a similar manner. Each period, the amount of uncertainty is reduced by a constant factor $1 - \theta$, such that the updated variance of beliefs in period t is,

$$\sigma^2_t = \theta\sigma^2_{t-1}.$$

Then, the more weight that is placed on prior information, the slower the learning.

B.2 Tables

Table B.1 Description of the Pre-ITQ Fishery

	Percentile					**Total**
	10	**25**	**50**	**75**	**90**	
Fisherman skill (d)	0.0505	0.1278	0.2910	0.5350	0.7717	-
Harvest (q_0)	0.1791	0.2011	0.2305	0.2556	0.2695	10.0000
Profits (π_0)	0.0104	0.0242	0.0427	0.0584	0.0671	1.7890

Table B.2 Description of Transition

	Quota Price (r_t)	**Per Period Price**	**Fleet Size** (n_t)
Pre-ITQ	-	-	44.00
$t = 1$	0.6836	0.2522	39.60
$t = 2$	0.4492	0.2364	33.44
$t = 3$	0.2312	0.2312	30.92

Table B.3 Description of Fleet Configuration, Harvest, and Profits During Transition

	Percentile					**Total**
	10	**25**	**50**	**75**	**90**	
	Fisherman skill (d)					
Pre-ITQ	0.0505	0.1278	0.2910	0.5350	0.7717	-
$t = 1$	0.0458	0.1141	0.2539	0.4487	0.6174	-
$t = 2$	0.0392	0.0954	0.2060	0.3482	0.4590	-
$t = 3$	0.0365	0.0880	0.1877	0.3124	0.4055	-
	Harvest (q)					
Pre-ITQ	0.1791	0.2011	0.2305	0.2556	0.2695	10.0000
$t = 1$	0.1956	0.2193	0.2548	0.2848	0.3054	10.0000
$t = 2$	0.2458	0.2660	0.2971	0.3272	0.3465	10.0000
$t = 3$	0.2735	0.2929	0.3237	0.3533	0.3709	10.0000
	Harvest profits (gross quota purchase)					
Pre-ITQ	0.0104	0.0242	0.0427	0.0584	0.0671	1.7890
$t = 1$	0.0341	0.0491	0.0732	0.0941	0.1073	2.8739
$t = 2$	0.0577	0.0706	0.0907	0.1098	0.1209	2.9810
$t = 3$	0.0684	0.0803	0.0992	0.1175	0.1283	3.0579
	Profits (net quota purchase)					
Pre-ITQ	0.0104	0.0242	0.0427	0.0584	0.0671	1.7890
$t = 1$	0.0318	0.0441	0.0606	0.0759	0.0846	2.8739
$t = 2$	0.0469	0.0595	0.0783	0.0939	0.1035	2.9810
$t = 3$	0.0643	0.0759	0.0940	0.1118	0.1222	3.0579

Table B.4 Rent Generation

	Sources of Rent			**Total**
	Skill	**Resource**	**Speculation**	
Pre-ITQ	1.7890	0.0000	-	1.7890
$t = 1$	0.3514	1.3343	1.1882	2.8739
$t = 2$	0.6167	1.4701	0.8943	2.9810
$t = 3$	0.7459	2.3120	0.0000	3.0579

Table B.5 Comparative Dynamics of Key Model Parameters

Marginal Change	Transition			Long Run Equilibrium
	Price Path	Vessel Exit	Fleet Configuration	
Mean Beliefs (μ)	Decreases early prices	Quickens the rate of vessel exit	More efficient fleet during transition	Unchanged
Entry cost (ψ)	Increases early prices	Slows the rate of vessel exit	Less efficient fleet during transition	Unchanged
Fixed cost (f)	Downward shift of the price path	Downward shift of exit path	Improves the efficiency of the fleet	Trading price lower, fleet size smaller, more efficient fleet
Time preference (δ)	Increases early prices	Slows the rate of vessel exit	Less efficient fleet during transition	Unchanged
Price of fish (p)	Upward shift of the price path	No effect on path	No effect on configuration	Trading price higher
Quota supply (Q)	Downward shift of the price path	Upward shift of exit path, Slows the rate of vessel exit	Less efficient during transition	Trading price lower, fleet size larger, less efficient fleet

Table B.6 Average Mean Belief of Fishermen on Transition Path

Mean Belief	**Ownership Price** $t=1$	$t=2$	$t=3$	**Per Period Price** $t=1$	$t=2$	$t=3$	**Fleet Size** *Pre-ITQ*	$t=1$	$t=2$	$t=3$
0.25	0.6887	0.4698	0.2309	0.2541	0.2473	0.2309	44.0	32.8	32.2	30.9
0.50	0.6879	0.4632	0.2308	0.2538	0.2438	0.2308	44.0	33.7	32.4	30.8
0.75	0.6883	0.4571	0.2312	0.2540	0.2406	0.2312	44.0	36.2	33.0	30.9
1.00	0.6836	0.4492	0.2312	0.2523	0.2364	0.2312	44.0	39.6	33.4	30.9
1.50	0.6317	0.4373	0.2314	0.2331	0.2302	0.2314	44.0	43.4	35.5	31.0
2.00	0.5565	0.4301	0.2316	0.2054	0.2264	0.2316	44.0	44.0	38.0	31.0
3.00	0.4378	0.4187	0.2316	0.1615	0.2204	0.2316	44.0	44.0	41.4	31.0

Table B.7 Entry Cost (ψ) on Transition Path

Entry Cost (ψ)	**Ownership Price** $t=1$	$t=2$	$t=3$	**Per Period Price** $t=1$	$t=2$	$t=3$	**Fleet Size** *Pre-ITQ*	$t=1$	$t=2$	$t=3$
0.00	0.4706	0.3967	0.2316	0.1737	0.2088	0.2316	44.0	30.4	30.6	31.0
0.05	0.6268	0.4348	0.2316	0.2313	0.2288	0.2316	44.0	37.2	32.7	31.0
0.10	0.6836	0.4492	0.2312	0.2522	0.2364	0.2312	44.0	39.6	33.4	30.9
0.15	0.7018	0.4549	0.2310	0.2590	0.2394	0.2310	44.0	40.0	34.0	30.9
0.20	0.7171	0.4609	0.2307	0.2646	0.2426	0.2307	44.0	40.4	34.4	30.8
0.25	0.7293	0.4660	0.2299	0.2691	0.2453	0.2299	44.0	40.4	34.8	30.8

Table B.8 Per Period Fixed Cost (f) on Transition Path

Fixed Cost (f)	**Ownership Price** $t=1$	$t=2$	$t=3$	**Per Period Price** $t=1$	$t=2$	$t=3$	**Fleet Size** *Pre-ITQ*	$t=1$	$t=2$	$t=3$
0.000	1.0478	0.8066	0.4479	0.3867	0.4246	0.4479	50.0	50.0	50.0	50.0
0.025	1.0455	0.8046	0.4479	0.3858	0.4235	0.4479	50.0	50.0	50.0	50.0
0.050	0.9682	0.7169	0.3881	0.3573	0.3773	0.3881	50.0	48.0	44.6	42.0
0.075	0.8308	0.5793	0.3045	0.3066	0.3049	0.3045	49.0	44.1	38.2	35.0
0.100	0.6836	0.4492	0.2312	0.2522	0.2364	0.2312	44.0	39.6	33.4	30.9
0.125	0.5630	0.3563	0.1721	0.2077	0.1876	0.1721	38.0	35.7	30.8	28.0

Table B.9 Time Preference (δ) on Transition Path

Time Pref. (δ)	Ownership Price			Per Period Price			Fleet Size			
	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	*Pre-ITQ*	$t=1$	$t=2$	$t=3$
0.00	0.2316	0.2316	0.2316	0.2316	0.2316	0.2316	44.0	31.0	31.0	31.0
0.01	0.2316	0.2316	0.2316	0.2293	0.2293	0.2316	44.0	31.0	31.0	31.0
0.10	0.2331	0.2334	0.2316	0.2100	0.2122	0.2316	44.0	31.0	31.0	31.0
0.25	0.2609	0.2501	0.2316	0.1988	0.2001	0.2316	44.0	31.0	31.0	31.0
0.50	0.3195	0.2931	0.2307	0.1826	0.1954	0.2307	44.0	32.6	31.2	30.9
0.75	0.4418	0.3548	0.2304	0.1911	0.2027	0.2304	44.0	35.2	32.2	30.8
0.90	0.5889	0.4139	0.2310	0.2173	0.2179	0.2310	44.0	38.1	33.1	30.9
0.99	0.6836	0.4492	0.2312	0.2302	0.2257	0.2312	44.0	39.6	33.4	30.9
1.00	0.7402	0.4693	0.2314	0.2467	0.2347	0.2314	44.0	39.9	33.9	31.0

Table B.10 Ex-Vessel Price for Fish (p) on Transition Path

Fish Price (p)	Ownership Price			Per Period Price			Fleet Size			
	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	*Pre-ITQ*	$t=1$	$t=2$	$t=3$
1.0	0.6836	0.4492	0.2312	0.2522	0.2364	0.2312	44.0	39.6	33.4	31.0
1.1	0.9405	0.6290	0.3316	0.3470	0.3311	0.3316	47.0	40.0	33.1	31.0
1.2	1.2050	0.8161	0.4316	0.4446	0.4295	0.4316	49.0	40.4	32.9	31.0

Table B.11 Quota Supply (Q) on Transition Path

Quota Supply (Q)	Ownership Price			Per Period Price			Fleet Size			
	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$	*Pre-ITQ*	$t=1$	$t=2$	$t=3$
5.000	0.9711	0.5906	0.3099	0.3583	0.3108	0.3099	33.0	21.6	16.0	15.0
7.500	0.8458	0.5275	0.2733	0.3121	0.2776	0.2733	40.0	31.8	24.8	23.0
10.000	0.6836	0.4492	0.2312	0.2522	0.2364	0.2312	44.0	39.6	33.4	30.9
12.500	0.4750	0.3494	0.1791	0.1753	0.1839	0.1791	46.0	43.7	41.5	38.0
15.000	0.2450	0.2060	0.1071	0.0904	0.1084	0.1071	46.0	46.0	45.9	43.0

B.3 Figures

Figure B.1 Timing of Events

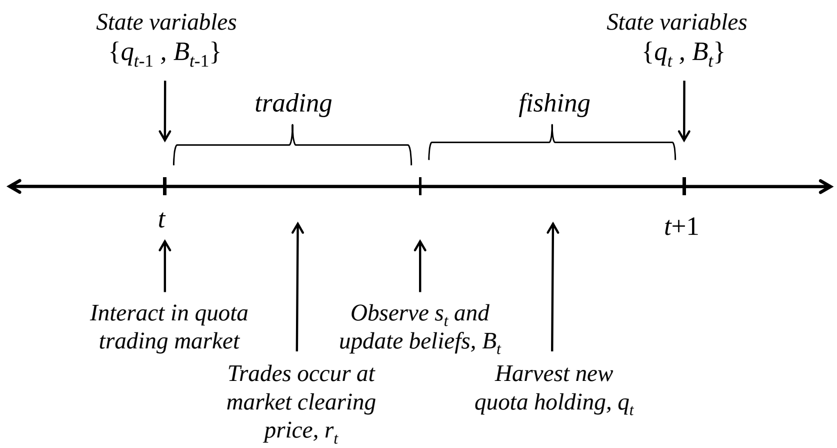

Figure B.2 Quota Holding During Transition

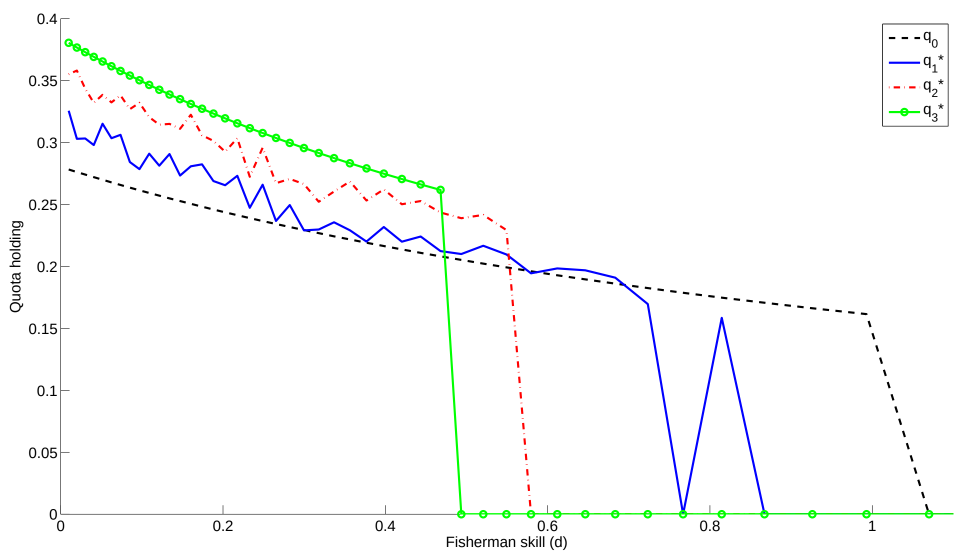

Figure B.3 Average Mean Belief of Fishermen on Transition

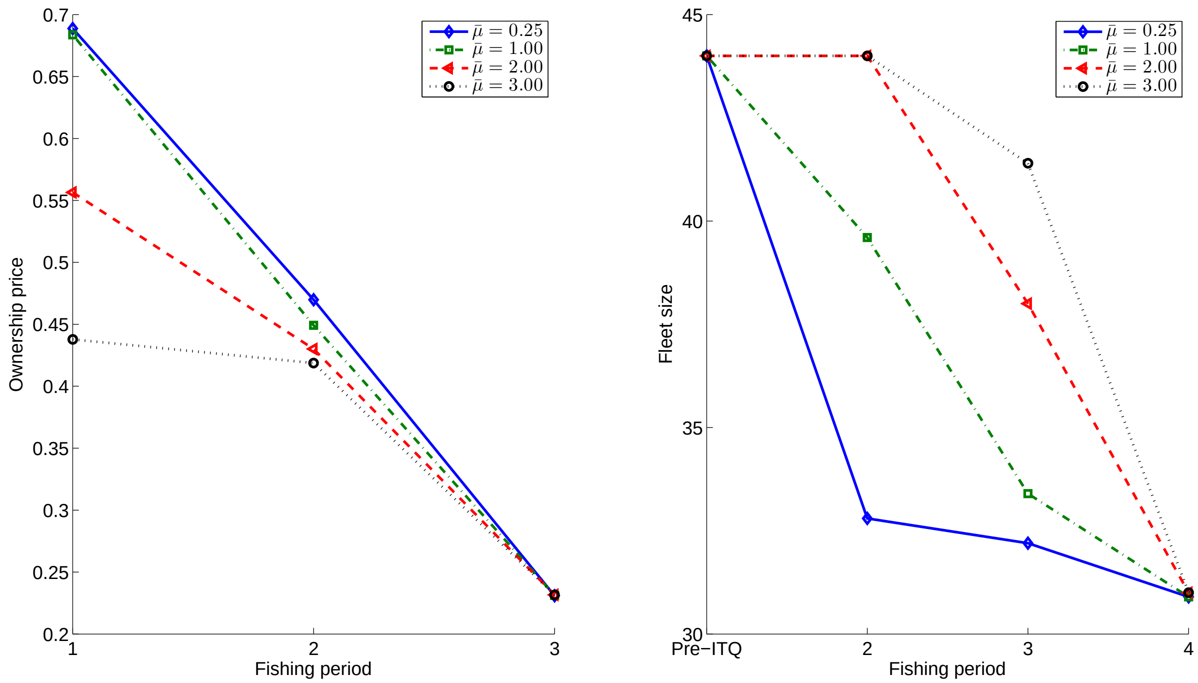

Figure B.4 Entry Cost (ψ) on Transition

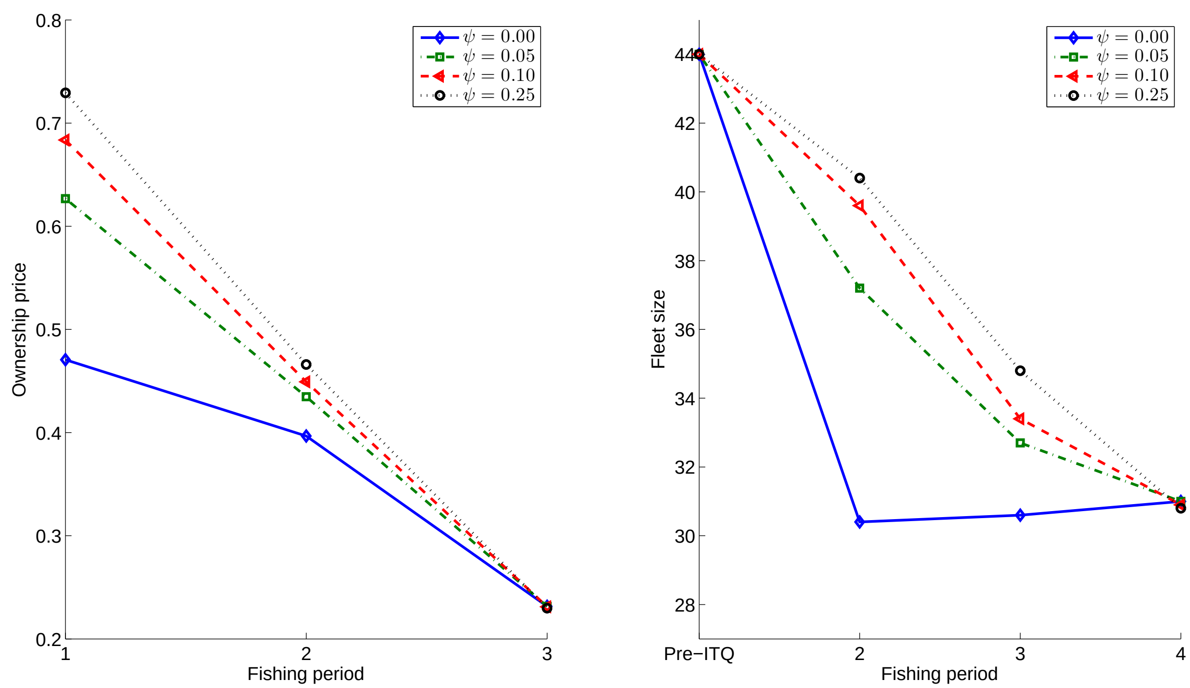

Figure B.5 Per Period Fixed cost (f) on Transition

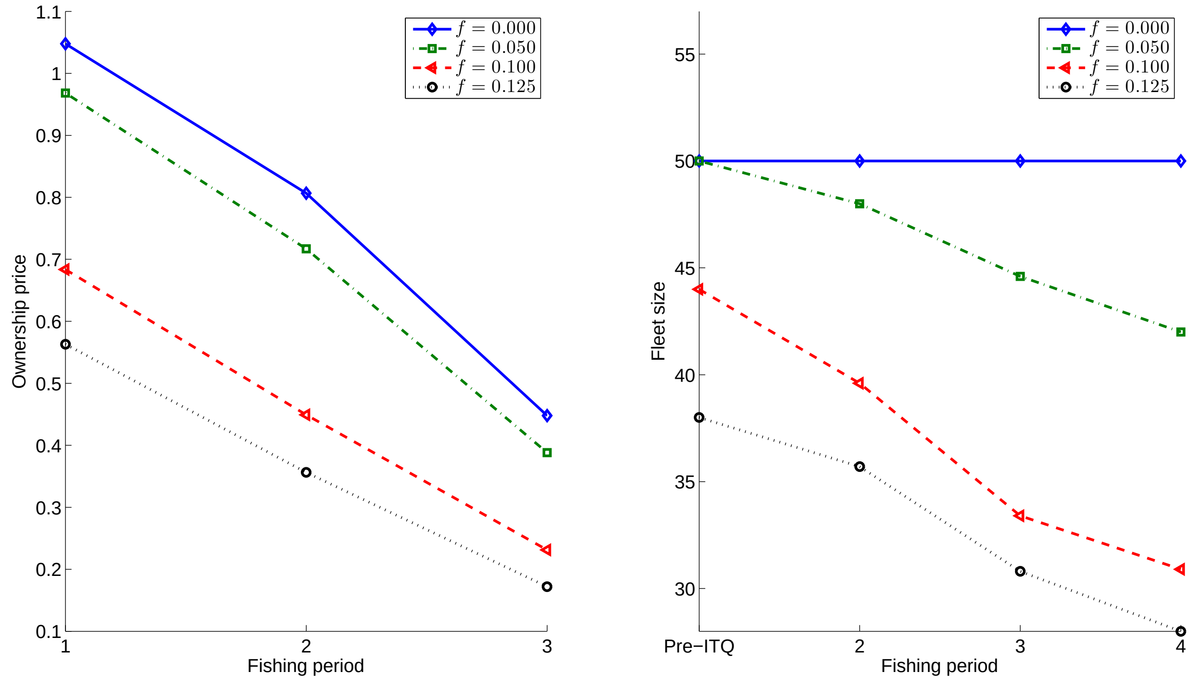

Figure B.6 Time Preference (δ) on Transition

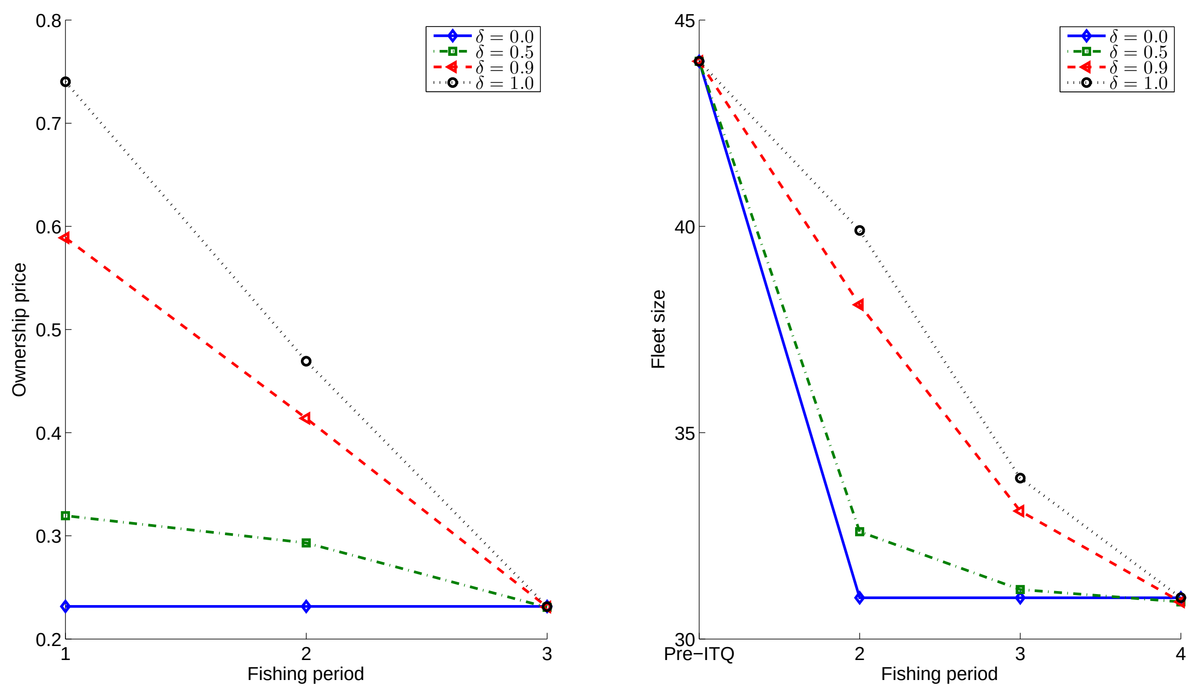

Figure B.7 Ex-Vessel Price for Fish (p) on Transition

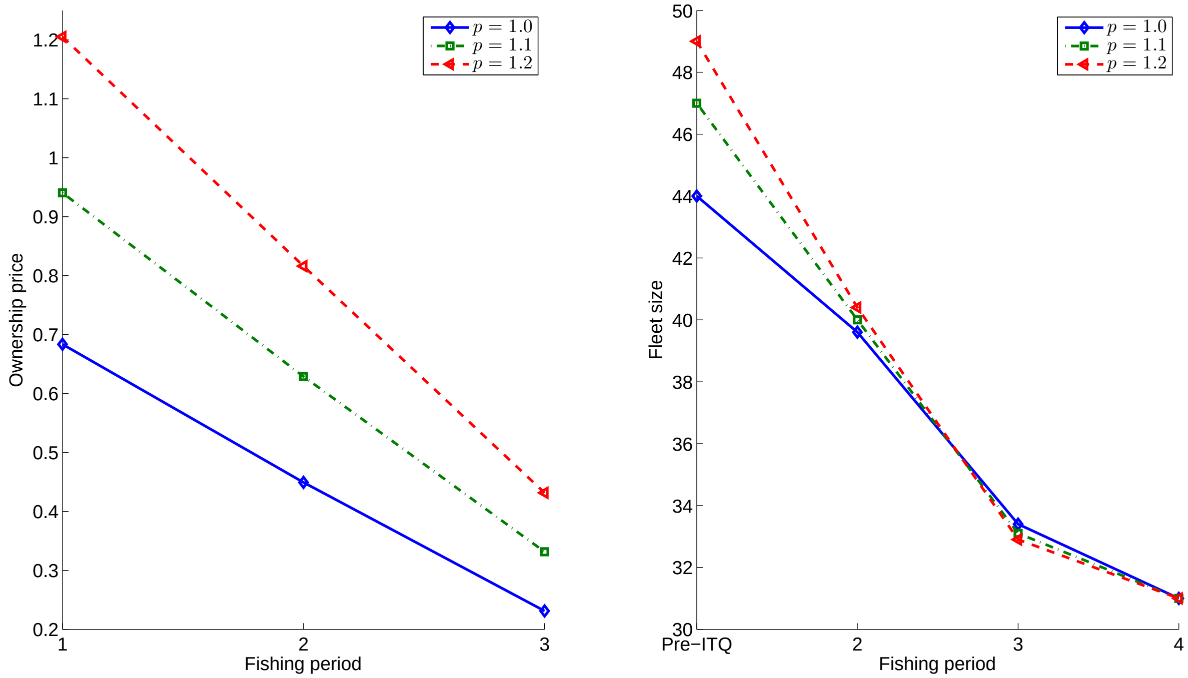

Figure B.8 Quota Supply (Q) on Transition

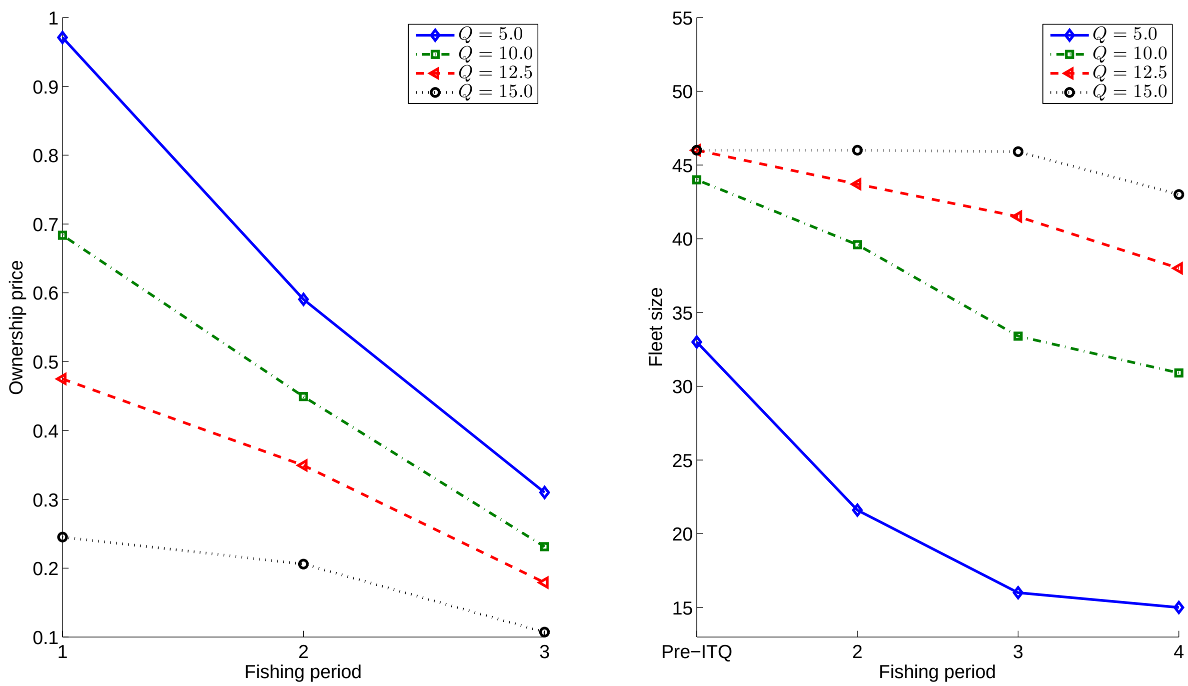

APPENDIX C. ADDITIONAL MATERIAL FOR CHAPTER 4

C.1 Tables

Table C.1 Timeline for Remediation of Superfund Sites in Maricopa County

Name (HRS score)	Proposed	Final Listing	Completion	Deleted
19th Avenue Landfill, Phoenix (54.27)	12/30/1982	09/08/1983	02/17/1998	09/25/2006
Phoenix-Goodyear Airport, Goodyear (45.91)	12/30/1982	09/08/1983	–	–
Indian Bend Wash Area, Scottsdale (42.24)	12/30/1982	09/08/1983	09/28/2006	05/01/2003*
Motorolla - Semiconductor, Phoenix (40.83)	10/15/1984	10/04/1989	–	–
Hassayampa Landfill, Hassayampa (42.79)	06/10/1986	07/22/1987	09/30/1997	–
Luke Air Force Base, Glendale (37.93)	07/14/1989	08/30/1990	09/25/2000	04/22/2002
Williams Air Force Base, Mesa (37.93)	07/14/1989	11/21/1989	–	–

* This site is partially deleted from the National Priority List.

Table C.2 User Cost for Imputing Annualized Rent

Variable	Model I	Model II
Interest rate (i)	0.0748	0.0672
Property tax rate (τ_p)	0.0154	0.0154
Risk preference (β)	0.0400	0.0400
Depreciation rate (δ)	0.0200	0.0200
Maintenance rate (m)	0.0200	0.0200
Inflation (π)	0.0342	0.0346
User cost	0.1360	0.1280

Table C.3 Summary Statistics of Housing Markets Data

School District	**Aguila Elem. Dist.**	**Saddle Mountain Unif. Dist.**	**Phoenix Union H.S. Dist.**	**Fountain Hills Unif. Dist.**	**Scottsdale Unif. Dist.**	**Full Sample**
			Mean			
Market Value ($)	78,600	83,000	104,989	260,226	293,943	162,792
Lot Acres	4.36	4.57	0.21	0.33	0.31	0.24
Square Feet (100s)	18.64	13.32	14.15	22.61	23.62	18.81
Floors	1.00	1.00	1.03	1.08	1.12	1.15
Bathrooms	2.42	1.57	1.81	3.07	3.05	2.60
Age (years)	11.83	29.67	38.37	7.24	18.46	14.41
Garage (*proportion*)	0.83	0.67	0.81	1.00	0.97	0.96
Pool (*proportion*)	0.00	0.33	0.17	0.37	0.50	0.24
			Median			
Market Value ($)	47,300	65,000	81,830	240,000	239,045	134,000
Lot Acres	5.00	4.63	0.17	0.26	0.21	0.18
Square Feet (100s)	16.69	14.52	13.09	21.63	21.98	17.34
Floors	1.00	1.00	1.00	1.00	1.00	1.00
Bathrooms	2.50	1.00	2.00	2.70	2.70	2.50
Age	14.00	33.00	42.00	4.00	14.00	8.00

Table C.4 Fixed Effects Hedonic Regression

Variable	**Model I**			**Model II**		
	Estimate	**Std. Err.**	**T-stat**	**Estimate**	**Std. Err.**	**T-stat**
Lot Acres	0.2214	0.0036	61.5733	0.2396	0.0029	83.2994
Square Feet	0.0614	0.0006	110.4043	0.0555	0.0004	123.8480
Floors	-0.1688	0.0022	-75.0229	-0.1714	0.0017	-99.7336
Bathrooms	0.0932	0.0018	51.0877	0.0990	0.0015	67.6634
Age	-0.0118	0.0002	-73.4769	-0.0110	0.0001	-82.0344
Lot Acres Sq	-0.0136	0.0004	-33.3631	-0.0148	0.0003	-44.4649
Square Feet Sq	-0.0005	0.0000	-43.7288	-0.0004	0.0000	-45.8566
Age Sq	0.0001	0.0000	49.1352	0.0001	0.0000	62.6331
Garage	0.1303	0.0040	32.5354	0.1065	0.0035	30.5513
Pool	0.0875	0.0018	47.7517	0.0812	0.0017	49.0944
Year 1998	8.6574	0.0087	994.0957	–	–	–
Year 1999	8.7176	0.0087	1001.5492	–	–	–
Year 2000	8.7887	0.0088	1003.5424	8.7814	0.0074	1184.4289
Year 2001	–	–	–	8.8179	0.0073	1214.9692
Year 2002	–	–	–	8.8677	0.0073	1220.2983
Year 2003	–	–	–	8.9182	0.0074	1204.2004
District 400001	0.2222	0.0059	37.6375	0.3386	0.0050	67.8123
District 400450	-0.1102	0.0062	-17.8401	-0.0542	0.0044	-12.2448
District 400480	-1.5116	0.0978	-15.4606	-0.3622	0.0667	-5.4336
District 401410	-0.1950	0.0103	-18.8586	-0.0822	0.0079	-10.3940
District 401870	-0.0198	0.0049	-4.0718	0.0707	0.0039	17.9642
District 402690	-0.0692	0.0053	-13.1281	-0.0445	0.0041	-10.8324
District 403040	0.2245	0.0077	29.0444	0.3854	0.0072	53.5390
District 403310	-0.4991	0.0548	-9.1137	-0.5038	0.0400	-12.5847
District 403400	-0.0635	0.0048	-13.1087	0.0286	0.0039	7.2602
District 403450	-0.1069	0.0048	-22.4887	0.0043	0.0039	1.1013
District 403780	-0.0833	0.0089	-9.3379	-0.0051	0.0049	-1.0502
District 404970	-0.0490	0.0045	-11.0039	0.0527	0.0037	14.4067
District 405340	-0.2306	0.0489	-4.7201	-0.2253	0.0452	-4.9863
District 405460	-0.5631	0.0663	-8.4961	-0.3705	0.0254	-14.6011
District 405930	0.0773	0.0047	16.4961	0.2240	0.0039	57.3629
District 406250	-0.1087	0.0048	-22.8009	-0.0042	0.0040	-1.0426
District 406330	-0.1025	0.0047	-21.6258	0.0403	0.0039	10.4644
District 406810	-0.3496	0.0186	-18.7674	-0.1342	0.0098	-13.7598
District 407170	-0.8475	0.1379	-6.1469	-1.2623	0.1029	-12.2713
District 407570	0.3035	0.0049	61.8265	0.4302	0.0041	104.0446
District 407750	-0.0624	0.0048	-12.9528	0.0713	0.0039	18.2698
District 408340	0.0851	0.0047	18.0077	0.1964	0.0040	49.6808
District 408520	-0.1966	0.0057	-34.7093	-0.0912	0.0043	-21.4303
District 409190	-0.0868	0.0168	-5.1633	0.0488	0.0149	3.2827

Table C.5 Rank Ordering of Communities (by Community Price)

District ID	**District Name**	**Model I Rank**	**Model II Rank**	**Change in Rank**
District 400001	Cave Creek Unified District	22	22	0
District 400450	Agua Fria Union High School District	9	9	0
District 400480	Aguila Elementary District	1	4	3
District 401410	Buckeye Union High School District	8	8	0
District 401870	Chandler Unified District	19	18	-1
District 402690	Dysart Unified District	15	10	-5
District 403040	Fountain Hills Unified District	23	23	0
District 403310	Gila Bend Unified District	4	2	-2
District 403400	Gilbert Unified District	16	14	-2
District 403450	Glendale Union High School District	11	13	2
District 403780	Higley Elementary School District	14	11	-3
District 404970	Mesa Unified School District	18	17	-1
District 405340	Morristown Elementary District	6	5	-1
District 405460	Nadaburg Elementary District	3	3	0
District 405930	Paradise Valley Unified District	20	21	1
District 406250	Peoria Unified District	10	12	2
District 406330	Phoenix Union High School District	12	15	3
District 406810	Queen Creek Unified District	5	6	1
District 407170	Saddle Mountain Unified District	2	1	-1
District 407570	Scottsdale Unified District	24	24	0
District 407750	Deer Valley Unified District	17	19	2
District 408340	Tempe Union High School District	21	20	-1
District 408520	Tolleson Union High School District	7	7	0
District 409190	Wickenburg Unified District	13	16	3

School districts are labeled from 1 to 24, denoting the lowest to highest quality districts.

Table C.6 Fixed Effects Hedonic Regression - Anticipation

Variable	**Model I**			**Model II**		
	Estimate	**Std. Err.**	**T-stat**	**Estimate**	**Std. Err.**	**T-stat**
Lot Acres	0.2169	0.0041	52.7786	0.2246	0.0033	68.0954
Square Feet	0.0620	0.0006	97.1957	0.0552	0.0005	107.4580
Floors	-0.1641	0.0026	-63.2753	-0.1725	0.0020	-87.4291
Bathrooms	0.0920	0.0021	43.8952	0.0984	0.0017	58.6468
Age	-0.0117	0.0002	-63.6210	-0.0110	0.0002	-71.7413
Lot Acres Sq	-0.0132	0.0005	-28.7097	-0.0136	0.0004	-36.2616
Square Feet Sq	-0.0005	0.0000	-39.5900	-0.0004	0.0000	-39.1691
Age Sq	0.0001	0.0000	41.2237	0.0001	0.0000	54.5129
Garage	0.1343	0.0046	29.4760	0.1047	0.0040	26.3070
Pool	0.0857	0.0021	41.4832	0.0824	0.0019	43.3848
Year 1998	8.7218	0.0099	877.5295	–	–	–
Year 1999	8.7817	0.0099	883.6209	–	–	–
Year 2000	8.8308	0.0102	864.9584	8.8477	0.0085	1044.7155
Year 2001	–	–	–	8.8841	0.0083	1065.7049
Year 2002	–	–	–	8.9299	0.0084	1068.2867
Year 2003	–	–	–	–	–	–
District 400001	0.2080	0.0067	30.8368	0.3454	0.0058	59.6714
District 400450	-0.1205	0.0072	-16.8472	-0.0442	0.0052	-8.5296
District 400480	-1.5047	0.0988	-15.2308	-0.2487	0.0786	-3.1642
District 401410	-0.2161	0.0122	-17.6700	-0.0655	0.0091	-7.1671
District 401870	-0.0263	0.0055	-4.7649	0.0775	0.0046	16.8570
District 402690	-0.0569	0.0061	-9.3402	-0.0387	0.0047	-8.1584
District 403040	0.2025	0.0088	23.1007	0.3899	0.0083	46.8920
District 403310	-0.4448	0.0668	-6.6626	-0.4869	0.0436	-11.1679
District 403400	-0.0717	0.0055	-12.9987	0.0338	0.0046	7.3845
District 403450	-0.1187	0.0054	-22.0054	0.0103	0.0046	2.2545
District 403780	-0.0870	0.0111	-7.8226	0.0043	0.0057	0.7550
District 404970	-0.0580	0.0051	-11.4880	0.0594	0.0043	13.9390
District 405340	-0.1252	0.0554	-2.2595	-0.1823	0.0514	-3.5486
District 405460	-0.5092	0.0910	-5.5964	-0.3396	0.0295	-11.5219
District 405930	0.0606	0.0053	11.4222	0.2247	0.0046	49.3562
District 406250	-0.1187	0.0054	-21.9858	-0.0025	0.0046	-0.5461
District 406330	-0.1204	0.0054	-22.2921	0.0414	0.0045	9.2448
District 406810	-0.3986	0.0228	-17.4967	-0.1519	0.0121	-12.5480
District 407170	-0.7594	0.1701	-4.4636	-1.5183	0.1175	-12.9234
District 407570	0.2874	0.0056	51.7131	0.4301	0.0048	89.2129
District 407750	-0.0787	0.0055	-14.3726	0.0756	0.0045	16.6471
District 408340	0.0757	0.0054	14.1164	0.2015	0.0046	43.7905
District 408520	-0.2061	0.0065	-31.8728	-0.0846	0.0049	-17.1011
District 409190	-0.1030	0.0193	-5.3528	0.0749	0.0170	4.3962

Table C.7 Descriptive Statistics for the 24 School Districts

	Model I					Model II				
Variable	**Mean**	**Median**	**Std. Dev.**	**Min.**	**Max.**	**Mean**	**Median**	**Std. Dev.**	**Min.**	**Max.**
Community Size (population share)	0.042	0.015	0.054	0.0004	0.206	0.042	0.016	0.054	0	0.21
Community Price (relative)	3.999	4.125	1.145	1	6.142	3.516	3.534	0.97	1	5.433
AIMS Math (scaled mean)	679.667	698	60.885	505	726	–	–	–	–	–
AIMS Reading (scaled mean)	672.417	697	71.448	472	723	–	–	–	–	–
AIMS Writing (scaled mean)	668.375	696	66.508	449	717	–	–	–	–	–
Meet AIMS Standard (%)	60.444	63.667	11.108	27.667	71.667	–	–	–	–	–
Exceed AIMS Standard (%)	11.486	11.167	7.323	0	22.667	–	–	–	–	–
O_3 Exceedences (days)	18.423	17.645	10.927	2.675	46.849	10.138	8.042	6.191	2.675	24.069
O_3 Max (ppm)	0.086	0.087	0.003	0.079	0.09	0.091	0.089	0.006	0.085	0.102
O_3 90 (ppm)	0.072	0.072	0.002	0.067	0.075	0.07	0.069	0.001	0.067	0.073
O_3 Mean (ppm)	0.052	0.053	0.006	0.04	0.058	0.053	0.053	0.004	0.043	0.058
PM_{10} Max ($\mu g/m^3$)	112.09	109.042	16.65	78.847	146.185	127.705	128.775	16.814	96.681	152.752
PM_{10} 90 ($\mu g/m^3$)	73.573	76.245	10.271	51.673	90.398	65.707	66.588	9.767	49.457	83.308
PM_{10} Mean ($\mu g/m^3$)	44.167	45.699	6.801	28.754	55.551	42.901	43.682	6.08	31.997	53.483
Nearest NPL site (miles)	15.308	10.182	12.635	3.256	53.523	17.166	12.424	13.869	3.919	57.433
Income 25 (dollars)	31051	31070	9536	17500	52479	31135	31762	9622	17500	53380
Income 50 (dollars)	50722	49224	14242	32470	88756	50737	48904	14255	32470	88756
Income 75 (dollars)	78829	71648	20470	53780	138450	78884	70249	21223	53228	140252
Housing 25 (dollars)	105114	109250	42898	28000	205000	126030	120083	42786	47750	244090
Housing 50 (dollars)	134510	139850	57255	47300	270475	161931	162500	58694	69900	322750
Housing 75 (dollars)	183453	170140	75661	84463	374900	215537	197606	87382	86000	437500

Data on school quality is collected for a single year.

Table C.8 Parameter Estimates

Variable	**Model I**			**Model II**			**Model III**		
	Estimate	**Std Err**	**t-stat**	**Estimate**	**Std Err**	**t-stat**	**Estimate**	**Std Err**	**t-stat**
Mean log income ($\mu_{\ln y}$)	10.803	0.015	697.007	10.809	0.013	824.676	10.883	0.042	258.753
Var log income ($\sigma^2_{\ln y}$)	0.381	-0.008	-46.348	0.370	-0.010	-38.837	0.362	0.016	22.440
Mean log taste ($\mu_{\ln \alpha}$)	0.038	0.004	9.603	0.059	0.006	10.021	0.051	0.007	6.762
Var log taste ($\sigma^2_{\ln \alpha}$)	0.063	-0.001	-54.728	0.084	-0.002	-41.583	0.071	0.007	9.997
Correlation (λ)	-0.326	0.005	-71.740	-0.320	0.004	-75.820	-0.309	0.023	-13.201
Public good in Comm. 1 (g_1)	0.472	0.021	22.008	0.292	0.015	20.136	–	–	–
Public good in Comm. 1 (g_1) (*before*)	–	–	–	–	–	–	0.362	0.018	19.568
Public good in Comm. 1 (g_1) (*after*)	–	–	–	–	–	–	0.374	0.017	22.644
Air quality (γ_2)	-0.011	0.001	-7.928	-0.040	0.013	-2.961	-0.025	0.003	-8.378
Inv. Dist. to nearest NPL site (γ_3)	-0.003	0.001	-4.407	-0.006	0.001	-7.742	-0.005	0.001	-5.552
Urban dummy (γ_4)	0.003	0.000	6.087	0.005	0.001	5.326	0.005	0.001	4.820
CES subst. parameter (ρ)	-0.019	0.000	-74.629	-0.023	0.000	-85.207	-0.021	0.001	-15.255
Demand scaler (β)	1.390	0.017	82.613	2.310	0.022	103.102	1.910	0.112	17.090
Price elasticity (η)	-0.231	0.016	-14.570	-0.584	0.023	-25.245	-0.411	0.038	-10.732
Income elasticity (ν)	0.783	0.002	328.886	0.795	0.003	253.290	0.774	0.010	78.478

The coefficient on air quality (γ_1) is fixed to 1.

Table C.9 One Standard Deviation Change in Public Good Component on Public Good Index

	Model I		Model II		Model III			
Variable	ΔX	Δg	ΔX	Δg	ΔX	Δg	ΔX	Δg
School quality	7.32	7.32	7.32	7.32	7.32	7.32	7.32	7.32
Air quality	6.8E-03	-7.8E-05	6.1E-03	-2.4E-04	6.8E-03	-1.7E-04	6.1E-03	-1.5E-04
Inv. dist. to nearest NPL site	6.9E-02	-2.3E-04	6.7E-02	-4.1E-04	6.9E-02	-3.6E-04	6.7E-02	-3.5E-04

ΔX denotes a one standard deviation change in public good attribute.

Table C.10 Inside Sample Fit

Fitted - Empirical

Model	Measure	1%	5%	15%	25%	35%	50%	65%	75%	85%	95%	99%	Mean
							Log Income						
Model I	Mean Deviation	0.74	0.45	0.09	-0.01	-0.07	-0.09	-0.09	-0.10	-0.10	-0.15	-0.32	-0.01
	Mean Absolute Deviation	0.74	0.48	0.21	0.20	0.20	0.20	0.19	0.18	0.18	0.20	0.32	0.23
Model II	Mean Deviation	0.76	0.46	0.09	-0.01	-0.06	-0.09	-0.10	-0.10	-0.10	-0.16	-0.33	-0.02
	Mean Absolute Deviation	0.76	0.49	0.21	0.20	0.21	0.20	0.19	0.18	0.18	0.20	0.33	0.23
							Log Housing Expenditures						
Model I	Mean Deviation	-0.11	-0.16	-0.10	-0.08	-0.04	0.01	0.05	0.07	0.08	0.07	0.01	-0.01
	Mean Absolute Deviation	0.26	0.25	0.16	0.16	0.15	0.16	0.18	0.19	0.19	0.19	0.21	0.18
Model II	Mean Deviation	-0.26	-0.27	-0.18	-0.11	-0.06	0.00	0.04	0.06	0.07	0.05	-0.02	-0.04
	Mean Absolute Deviation	0.42	0.30	0.21	0.15	0.14	0.16	0.17	0.18	0.20	0.20	0.21	0.19

(*Fitted - Empirical*)/*Empirical*

Model	Measure	1%	5%	15%	25%	35%	50%	65%	75%	85%	95%	99%	Mean	Mean*
							Log Income							*Income*
Model I	Mean Deviation (%)	8.54	4.98	1.02	-0.06	-0.57	-0.78	-0.82	-0.86	-0.81	-1.21	-2.54	0.02	3.47
	Mean Absolute Deviation (%)	8.54	5.33	2.17	1.96	1.88	1.87	1.74	1.60	1.52	1.65	2.54	2.17	24.47
Model II	Mean Deviation (%)	8.79	5.04	0.99	-0.04	-0.50	-0.78	-0.86	-0.88	-0.83	-1.29	-2.60	0.02	3.66
	Mean Absolute Deviation (%)	8.79	5.43	2.12	1.96	1.94	1.84	1.72	1.59	1.52	1.66	2.60	2.18	24.73
							Log Housing Expenditures							*Housing*
Model I	Mean Deviation (%)	-1.06	-1.67	-1.04	-0.81	-0.38	0.15	0.60	0.79	0.85	0.69	0.18	-0.07	1.87
	Mean Absolute Deviation (%)	3.00	2.63	1.65	1.60	1.54	1.59	1.84	1.91	1.89	1.81	1.91	1.82	17.99
Model II	Mean Deviation (%)	-2.69	-2.81	-1.84	-1.08	-0.59	0.01	0.46	0.67	0.74	0.54	-0.08	-0.39	-1.33
	Mean Absolute Deviation (%)	4.59	3.17	2.11	1.57	1.41	1.55	1.66	1.81	1.94	1.84	1.90	1.88	18.19

* denotes the mean associated with the unadjusted values of income and housing.
Mean values are calculated using quantiles 1 through 99.

Table C.11 Outside Sample Fit

Fitted - Empirical

Model	Measure	Percentile 1%	5%	15%	25%	35%	50%	65%	75%	85%	95%	99%	Mean
		Log Income											
Model I	Mean Deviation	0.73	0.44	0.08	-0.02	-0.07	-0.09	-0.10	-0.10	-0.09	-0.15	-0.32	-0.02
	Mean Absolute Deviation	0.73	0.48	0.21	0.20	0.20	0.20	0.19	0.18	0.17	0.20	0.32	0.22
Model II	Mean Deviation	0.75	0.47	0.10	-0.01	-0.06	-0.09	-0.10	-0.10	-0.10	-0.16	-0.33	-0.01
	Mean Absolute Deviation	0.75	0.49	0.21	0.19	0.19	0.19	0.18	0.17	0.17	0.20	0.33	0.22
		Log Housing Expenditures											
Model I	Mean Deviation	-0.32	-0.33	-0.24	-0.16	-0.12	-0.06	-0.01	0.01	0.01	-0.01	-0.08	-0.10
	Mean Absolute Deviation	0.45	0.34	0.25	0.19	0.17	0.16	0.16	0.16	0.17	0.16	0.19	0.19
Model II	Mean Deviation	-0.05	-0.11	-0.05	-0.03	0.02	0.07	0.11	0.13	0.14	0.13	0.07	0.05
	Mean Absolute Deviation	0.26	0.24	0.15	0.15	0.16	0.18	0.22	0.24	0.25	0.24	0.26	0.20

(Fitted - Empirical)/Empirical

Model	Measure	Percentile 1%	5%	15%	25%	35%	50%	65%	75%	85%	95%	99%	Mean	Mean*
		Log Income												*Income*
Model I	Mean Deviation (%)	8.49	4.84	0.87	-0.12	-0.56	-0.80	-0.85	-0.85	-0.78	-1.25	-2.52	-0.02	3.35
	Mean Absolute Deviation (%)	8.49	5.26	2.08	1.98	1.92	1.83	1.70	1.57	1.48	1.62	2.53	2.15	24.27
Model II	Mean Deviation (%)	8.76	5.15	1.11	-0.01	-0.53	-0.78	-0.83	-0.90	-0.85	-1.29	-2.65	0.05	3.58
	Mean Absolute Deviation (%)	8.76	5.44	2.16	1.87	1.81	1.78	1.65	1.52	1.44	1.63	2.65	2.12	24.01
		Log Housing Expenditures												*Housing*
Model I	Mean Deviation (%)	-3.33	-3.42	-2.42	-1.66	-1.16	-0.55	-0.10	0.11	0.19	-0.05	-0.65	-0.96	-6.79
	Mean Absolute Deviation (%)	4.85	3.61	2.53	1.91	1.68	1.60	1.61	1.61	1.60	1.44	1.71	1.93	17.78
Model II	Mean Deviation (%)	-0.42	-1.03	-0.44	-0.22	0.22	0.73	1.19	1.38	1.44	1.28	0.77	0.53	8.12
	Mean Absolute Deviation (%)	2.94	2.54	1.54	1.52	1.60	1.84	2.23	2.38	2.39	2.25	2.36	2.01	21.15

* denotes the mean associated with the unadjusted values of income and housing.
Mean values are calculated using quantiles 1 through 99.

Table C.12 WTP for Cleanup of Luke Air Force Base - Model I (*Ordered by transaction weighted distance from LAFB*)

Dist. Name (Rank)	$\mathbf{WTP}^{*}_{\mathbf{PE}}$	$\mathbf{WTP}^{*}_{\mathbf{GE}}$	**Share of (a)**	**Share of (b)**	$\mathbf{WTP_{GE}}$ (a)	(b)	(c)	$\mathbf{\Delta p}$ (a)	(b)	(c)	$\mathbf{\Delta g}$ (a)	(b)	(c)
Agua Fria Union H.S. Dist. (9)	0.00	1.22	1.8E-04	1.5E-04	0.92	3.24	1.22	0.006	0.330	-2.6E-04	0.047	2.489	0
Tolleson Union H.S. Dist. (7)	0.00	3.77	4.7E-05	1.9E-05	4.99	3.81	3.77	0.005	0.124	-8.1E-04	0.041	0.947	0
Dysart Unif. Dist. (15)	0.17	0.73	4.3E-04	3.9E-04	0.66	0.91	0.73	0.024	0.059	-1.1E-04	0.153	0.368	2.1E-04
Peoria Unif. Dist. (10)	0.09	1.22	5.0E-05	5.0E-05	1.16	0.92	1.22	0.007	0.006	-2.4E-04	0.057	0.047	1.4E-04
Glendale Union H.S. Dist. (11)	0.02	1.22	3.0E-05	2.5E-05	1.35	1.16	1.22	0.018	0.007	-2.4E-04	0.121	0.057	0
Buckeye Union H.S. Dist. (8)	0.00	3.82	2.7E-04	1.5E-04	3.24	4.99	3.82	0.330	0.005	-8.1E-04	2.489	0.041	0
Phoenix Union H.S. Dist. (12)	0.00	1.15	2.7E-05	1.6E-05	1.24	1.35	1.15	0.065	0.018	-2.3E-04	0.390	0.121	0
Deer Valley Unif. Dist. (17)	0.04	0.74	9.8E-05	1.0E-04	0.88	0.79	0.74	0.057	0.004	-1.3E-04	0.347	0.028	5.1E-05
Nadaburg Elem. Dist. (3)	0.08	24.50	7.3E-04	4.8E-04	19.70	14.55	24.50	-0.642	-0.176	-5.1E-03	-1.505	-0.567	5.0E-05
Paradise Valley Unif. Dist. (20)	0.00	-1.40	7.4E-05	8.1E-05	-1.42	0.08	-1.40	0.039	0.454	2.6E-04	0.245	2.816	0
Morristown Elem. Dist. (6)	0.02	6.13	9.5E-04	4.8E-04	3.81	17.44	6.13	0.124	0.403	-1.4E-03	0.947	2.474	2.6E-05
Tempe Union H.S. Dist. (21)	0.00	-1.53	4.1E-05	5.9E-05	-1.64	-1.42	-1.53	0.726	0.039	2.7E-04	4.716	0.245	0
Scottsdale Unif. Dist. (24)	0.00	-4.79	0	4.9E-05	–	-3.58	-4.79	–	0.467	8.3E-04	–	3.568	0
Saddle Mountain Unif. Dist. (2)	0.05	15.30	1.4E-03	1.0E-03	12.84	19.70	15.30	-0.933	-0.642	-3.0E-03	-1.041	-1.505	1.6E-05
Cave Creek Unif. Dist. (22)	0.00	-3.33	3.3E-04	3.5E-04	-3.47	-1.64	-3.33	0.014	0.726	6.0E-04	0.093	4.716	0
Chandler Unif. Dist. (19)	0.00	0.07	1.4E-04	1.7E-04	0.08	0.65	0.07	0.454	0.128	-1.3E-05	2.816	0.732	0
Fountain Hills Unif. Dist. (23)	0.00	-3.40	5.2E-04	5.5E-04	-3.58	-3.47	-3.40	0.467	0.014	6.1E-04	3.568	0.093	0
Mesa Unif. School Dist. (18)	0.00	0.68	5.1E-05	3.6E-05	0.65	0.88	0.68	0.128	0.057	-1.3E-04	0.732	0.347	0
Wickenburg Unif. Dist. (13)	0.01	1.00	2.3E-03	2.1E-03	0.95	1.24	1.00	0.014	0.065	-1.9E-04	0.089	0.390	1.2E-05
Gilbert Unif. Dist. (16)	0.00	0.72	1.2E-04	1.4E-04	0.79	0.66	0.72	0.004	0.024	-1.4E-04	0.028	0.153	0
Higley Elem. School Dist. (14)	0.00	0.95	3.5E-03	3.5E-03	0.91	0.95	0.95	0.059	0.014	-1.8E-04	0.368	0.089	0
Queen Creek Unif. Dist. (5)	0.00	15.54	1.1E-04	0	17.44	–	15.54	0.403	–	-3.3E-03	2.474	–	0
Gila Bend Unif. Dist. (4)	0.00	23.20	6.1E-04	0	14.55	–	23.20	-0.176	–	-4.9E-03	-0.567	–	0
Aguila Elem. Dist. (1)	0.04	-52.78	0	2.8E-03	–	12.84	-52.78	–	-0.933	9.2E-03	–	-1.041	4.3E-06

* Welfare measure is for the baseline population of school district j.
(**a**) denotes the households that move out of school district j.
(**b**) denotes the households that move into school district j.
(**c**) denotes the households that remain in school district j.

Table C.13 WTP for Discovery of Luke Air Force Base - Model II (*Ordered by transaction weighted distance from LAFB*)

Dist. Name (Rank)	$\mathbf{WTP^*_{PE}}$	$\mathbf{WTP^*_{GE}}$	**Share**	**Share**	$\mathbf{WTP_{GE}}$			$\mathbf{\Delta p}$			$\mathbf{\Delta g}$		
			of (a)	**of (b)**	(**a**)	(**b**)	(**c**)	(**a**)	(**b**)	(**c**)	(**a**)	(**b**)	(**c**)
Agua Fria Union H.S. Dist. (9)	0.00	6.30	7.1E-04	6.1E-04	6.43	7.26	6.30	0.031	0.092	-1.0E-03	0.277	0.792	0
Tolleson Union H.S. Dist. (7)	0.00	8.21	2.6E-04	5.4E-05	8.64	4.22	8.21	0.028	0.134	-1.3E-03	0.244	1.265	0
Dysart Unif. Dist. (10)	-0.28	5.79	9.9E-04	7.3E-04	5.27	6.43	5.79	0.135	0.031	-9.7E-04	1.185	0.277	-3.7E-04
Peoria Unif. Dist. (12)	-0.18	4.00	3.8E-04	2.5E-04	3.96	4.21	4.00	0.029	0.003	-6.8E-04	0.254	0.031	-2.5E-04
Glendale Union H.S. Dist. (13)	-0.04	3.58	2.8E-04	1.9E-04	3.56	3.96	3.58	0.087	0.029	-5.8E-04	0.680	0.254	0
Buckeye Union H.S. Dist. (8)	0.00	7.79	1.2E-03	8.5E-04	7.26	8.64	7.80	0.092	0.028	-1.2E-03	0.792	0.244	0
Phoenix Union H.S. Dist. (15)	0.00	1.88	1.8E-04	1.6E-04	1.99	2.22	1.88	0.031	0.042	-2.9E-04	0.210	0.330	0
Deer Valley Unif. Dist. (19)	-0.10	1.42	9.0E-04	9.2E-04	1.45	1.33	1.42	0.506	0.002	-2.3E-04	3.165	0.013	-9.2E-05
Nadaburg Elem. Dist. (3)	-0.13	-5.74	1.1E-03	1.1E-03	-5.00	3.97	-5.74	0.021	0.305	8.0E-04	0.097	1.265	-9.0E-05
Paradise Valley Unif. Dist. (21)	0.00	-3.12	4.8E-04	5.5E-04	-2.94	-1.83	-3.12	0.539	0.121	4.9E-04	3.266	0.751	0
Morristown Elem. Dist. (5)	-0.04	-1.51	1.3E-03	2.2E-03	-1.72	-4.74	-1.51	0.268	0.361	2.2E-04	2.16	2.117	-4.7E-05
Tempe Union H.S. Dist. (20)	0.00	-1.83	4.5E-04	5.1E-04	-1.83	1.45	-1.83	0.121	0.506	2.9E-04	0.751	3.165	0
Scottsdale Unif. Dist. (24)	0.00	-22.82	0	4.1E-04	–	-19.00	-22.82	–	0.242	3.8E-03	–	1.593	0
Saddle Mountain Unif. Dist. (1)	-0.47	40.68	1.4E-03	0	53.68	–	40.66	1.135	–	-3.6E-03	2.045	–	-2.8E-05
Cave Creek Unif. Dist. (22)	0.00	-13.58	2.5E-03	3.1E-03	-14.50	-2.94	-13.58	0.240	0.539	2.3E-03	1.534	3.266	0
Chandler Unif. Dist. (18)	0.00	1.42	1.1E-03	1.0E-03	1.33	1.55	1.42	0.002	0.067	-2.2E-04	0.013	0.419	0
Fountain Hills Unif. Dist. (23)	0.00	-18.27	4.2E-03	4.5E-03	-19.00	-14.50	-18.26	0.242	0.240	3.1E-03	1.593	1.534	0
Mesa Unif. School Dist. (17)	0.00	1.64	3.3E-04	3.0E-04	1.55	1.60	1.64	0.067	0.014	-2.5E-04	0.419	0.099	0
Wickenburg Unif. Dist. (16)	-0.02	1.71	1.6E-02	1.5E-02	1.60	1.99	1.71	0.014	0.031	-2.7E-04	0.099	0.210	-2.2E-05
Gilbert Unif. Dist. (14)	0.00	2.27	7.3E-04	6.6E-04	2.22	3.56	2.27	0.042	0.087	-3.6E-04	0.330	0.680	0
Higley Elem. School Dist. (11)	0.00	4.02	7.1E-03	8.3E-03	4.21	5.27	4.02	0.003	0.135	-6.6E-04	0.031	1.185	0
Queen Creek Unif. Dist. (6)	0.00	5.55	5.8E-04	2.9E-04	4.22	-1.72	5.55	0.134	0.268	-8.9E-04	1.265	2.16	0
Gila Bend Unif. Dist. (2)	0.00	3.50	1.4E-03	1.1E-03	3.97	53.68	3.49	0.305	1.135	-4.6E-04	1.265	2.045	0
Aguila Elem. Dist. (4)	-0.01	-5.98	3.2E-03	3.2E-03	-4.74	-5.00	-5.98	0.361	0.021	8.6E-04	2.12	0.097	-7.7E-06

* Welfare measure is for the baseline population of school district j.
(**a**) denotes the households that move out of school district j.
(**b**) denotes the households that move into school district j.
(**c**) denotes the households that remain in school district j.

C.2 Figures

Figure C.1 Household Sorting: Modified from Epple and Sieg (1999)

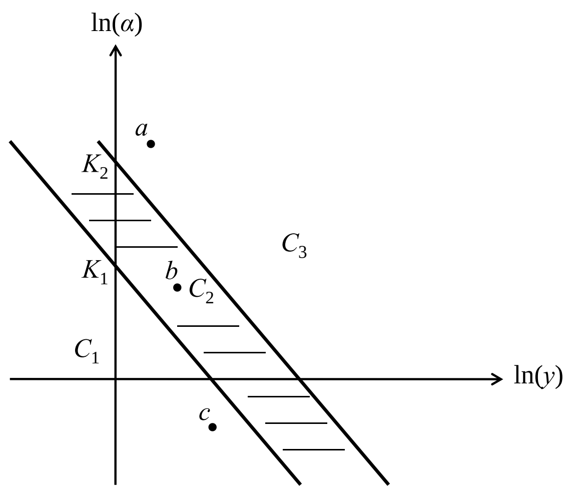

Figure C.2 School Districts and Superfund sites

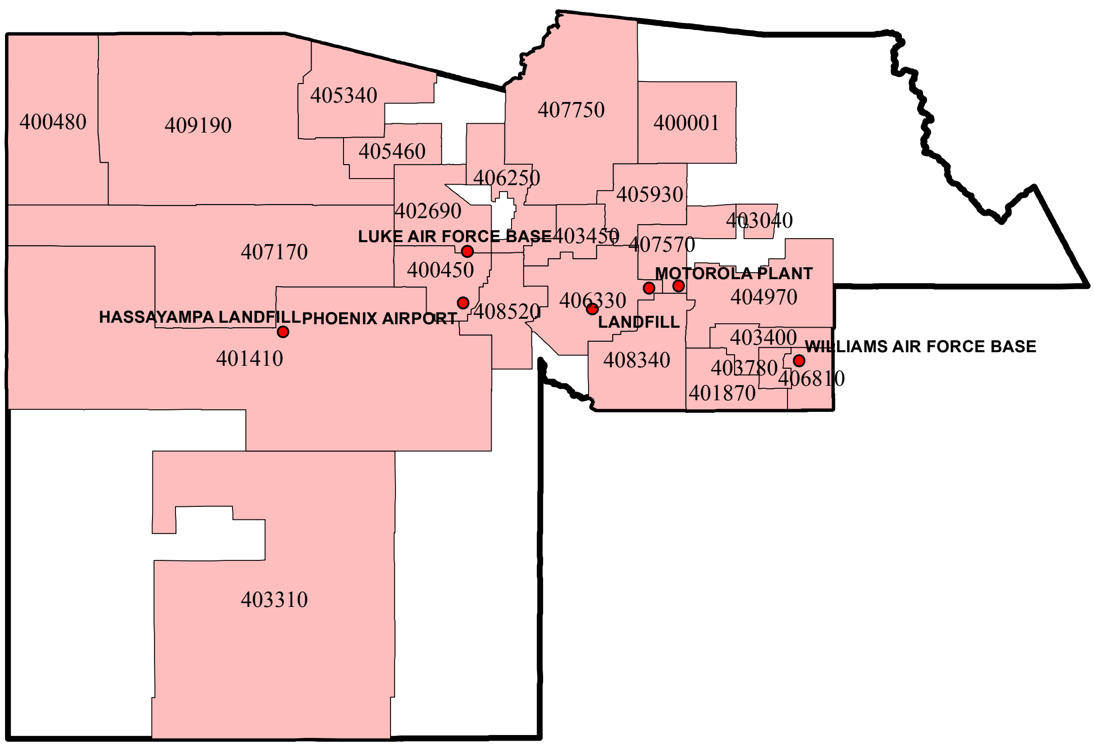

Figure C.3 Communities by Price Ranking - Model I

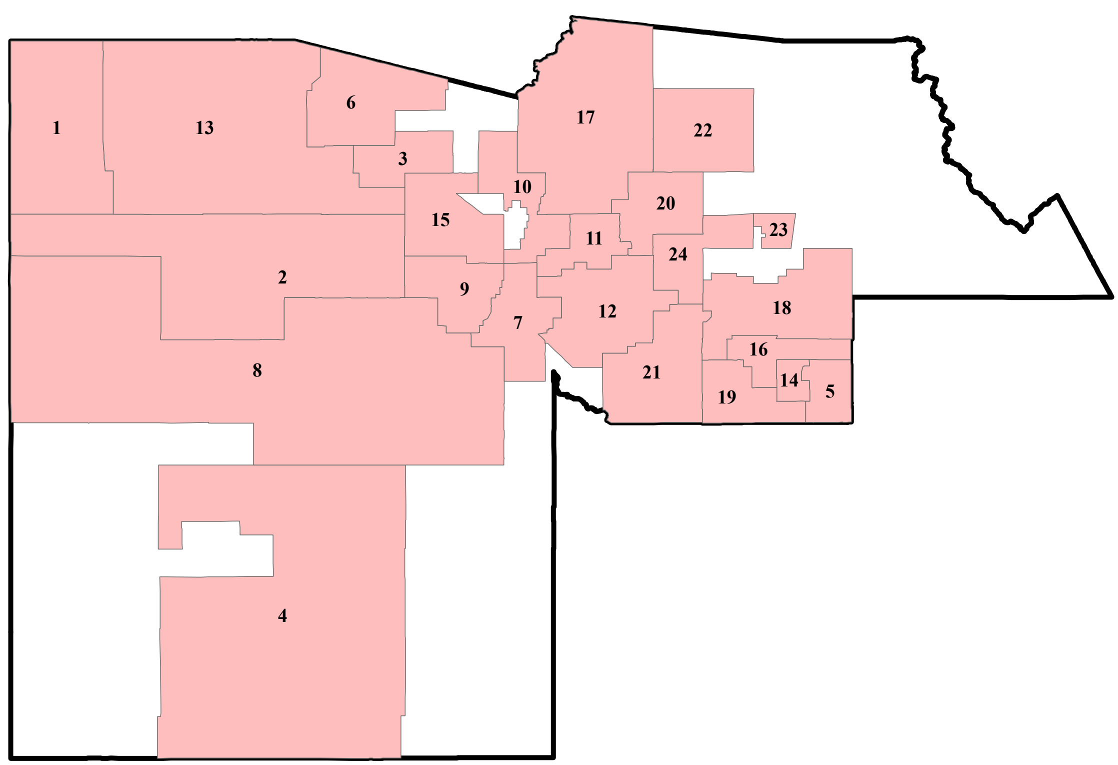

Figure C.4 Communities by Price Ranking - Model II

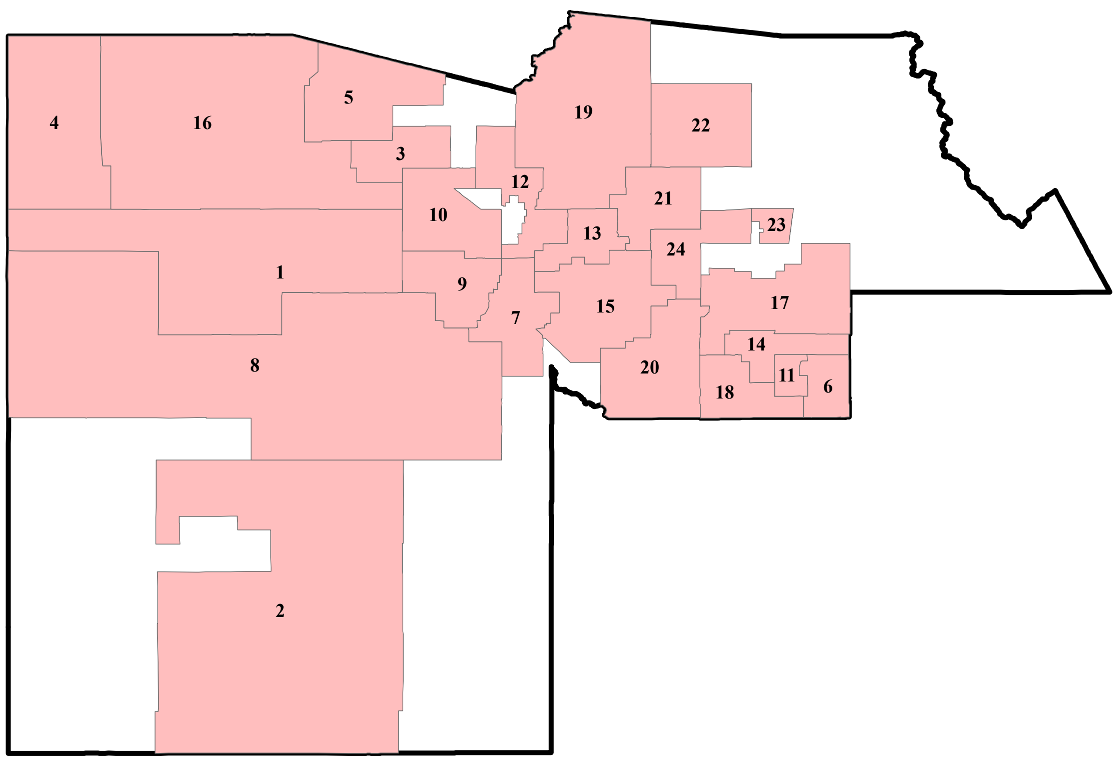

Figure C.5 Data Interpolation

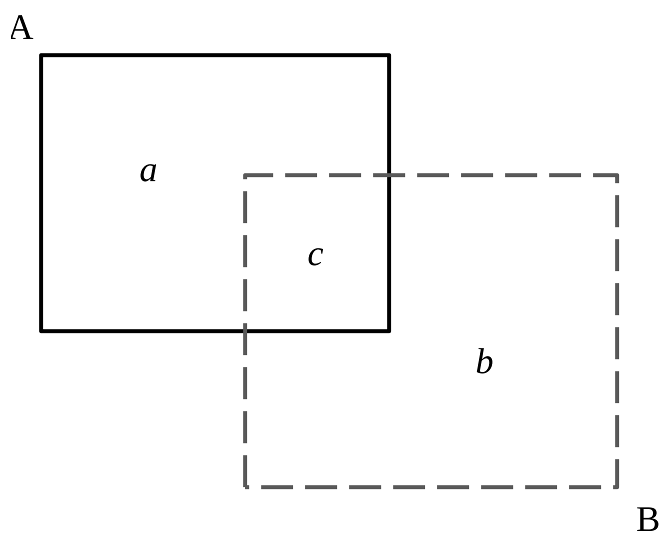

Figure C.6 Log Income - Model I

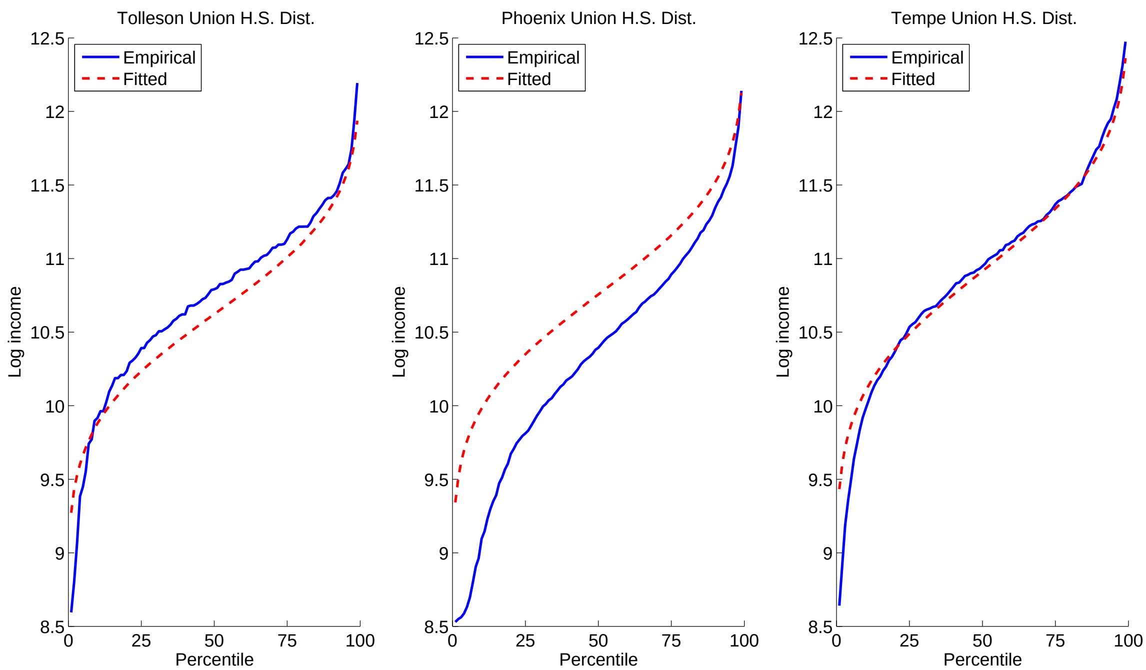

Figure C.7 Log Income - Model II

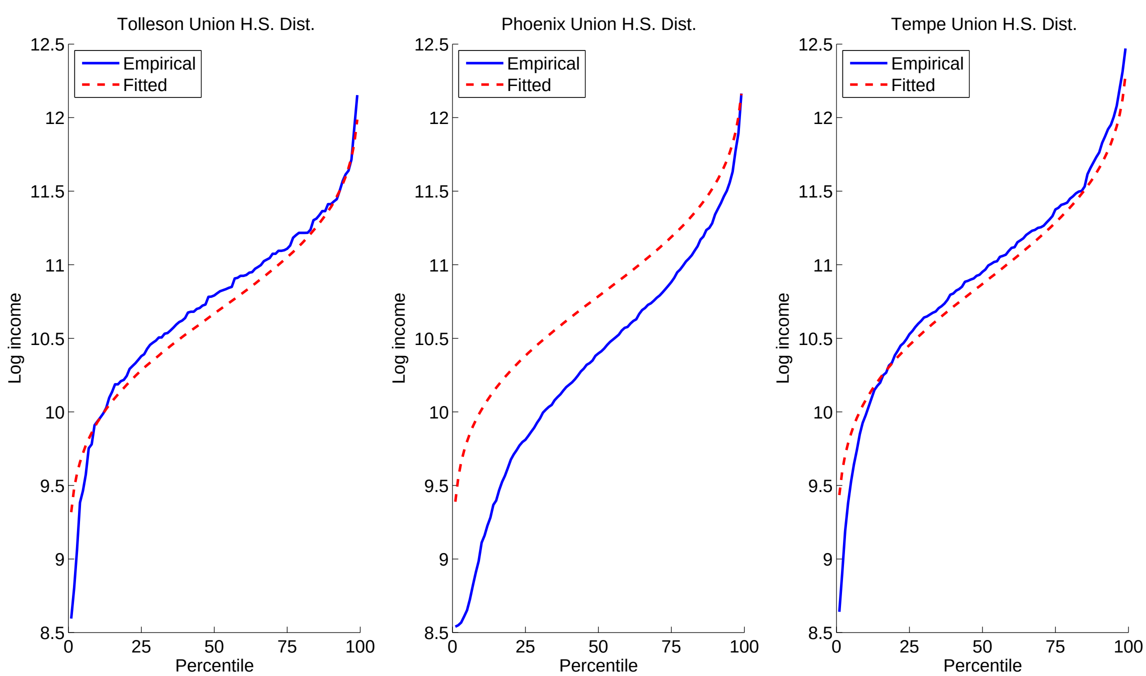

Figure C.8 Log Housing Expenditures - Model I

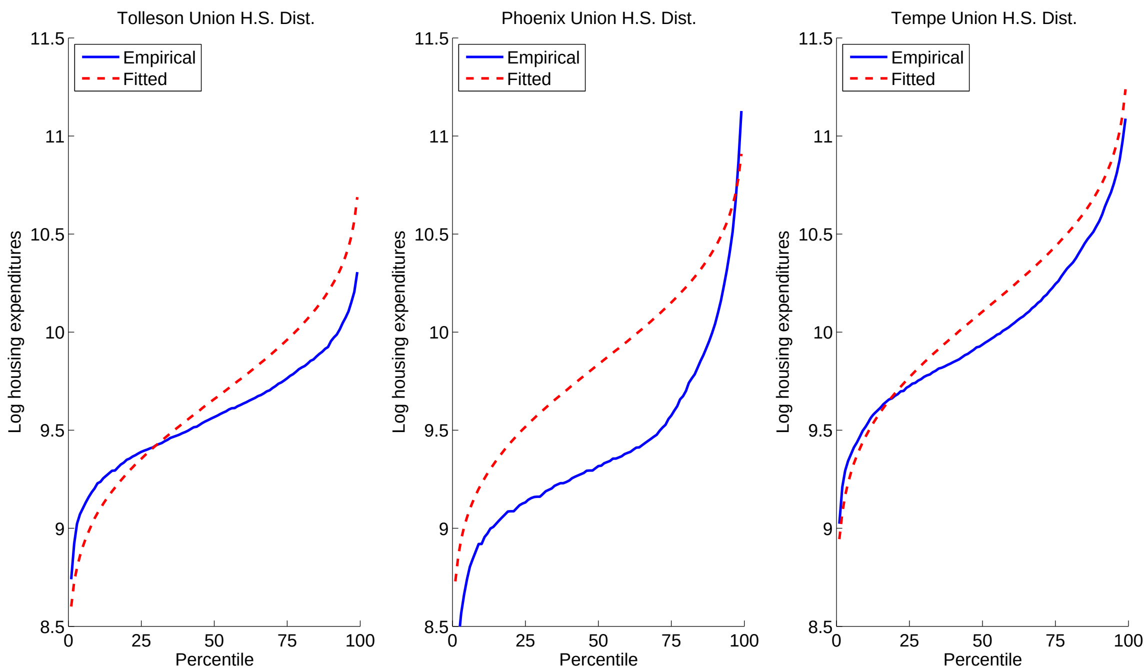

Figure C.9 Log Housing Expenditures - Model II

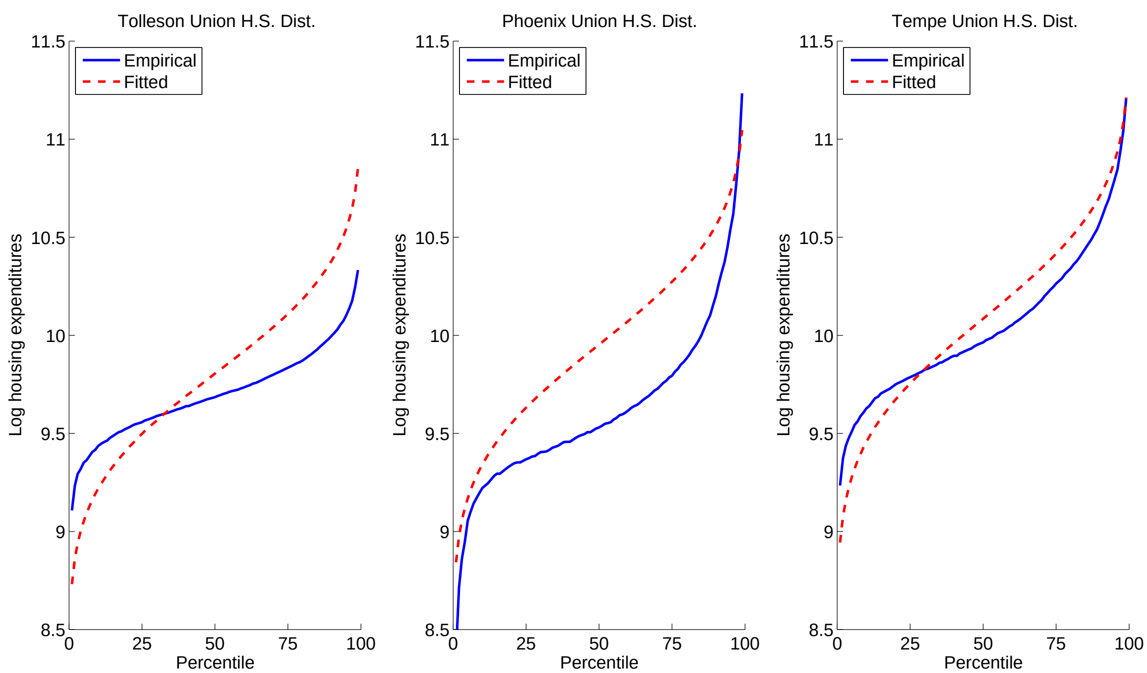

Figure C.10 Ascending Bundles - Model I

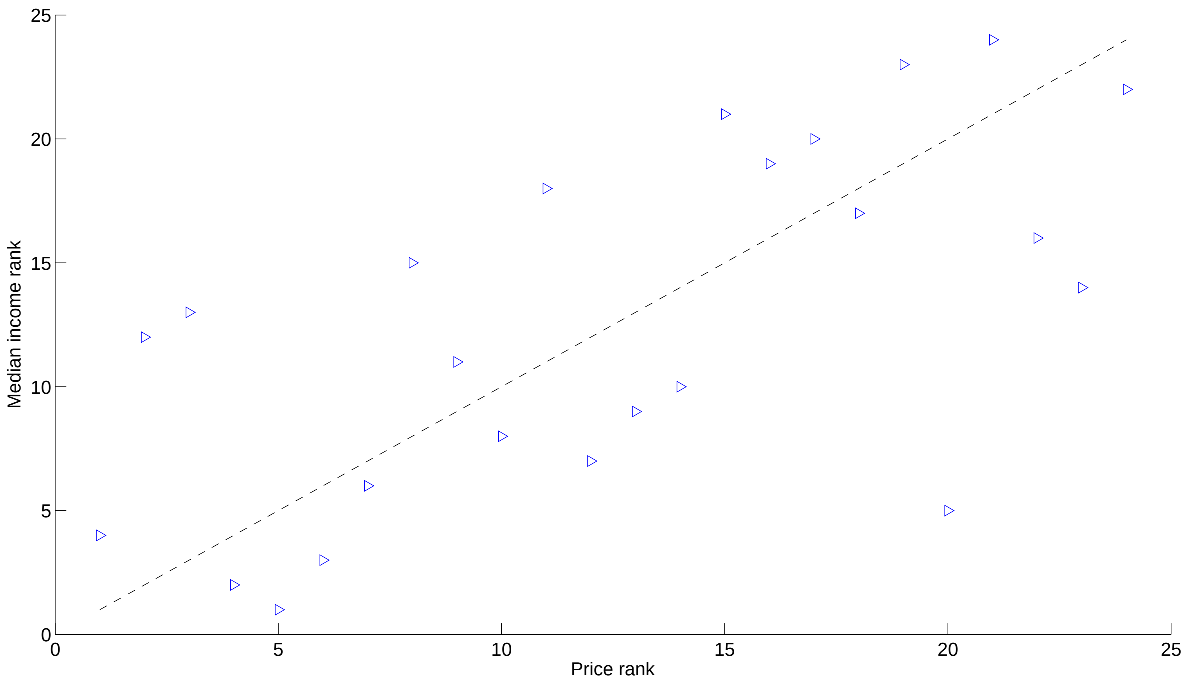

Figure C.11 Ascending Bundles - Model II

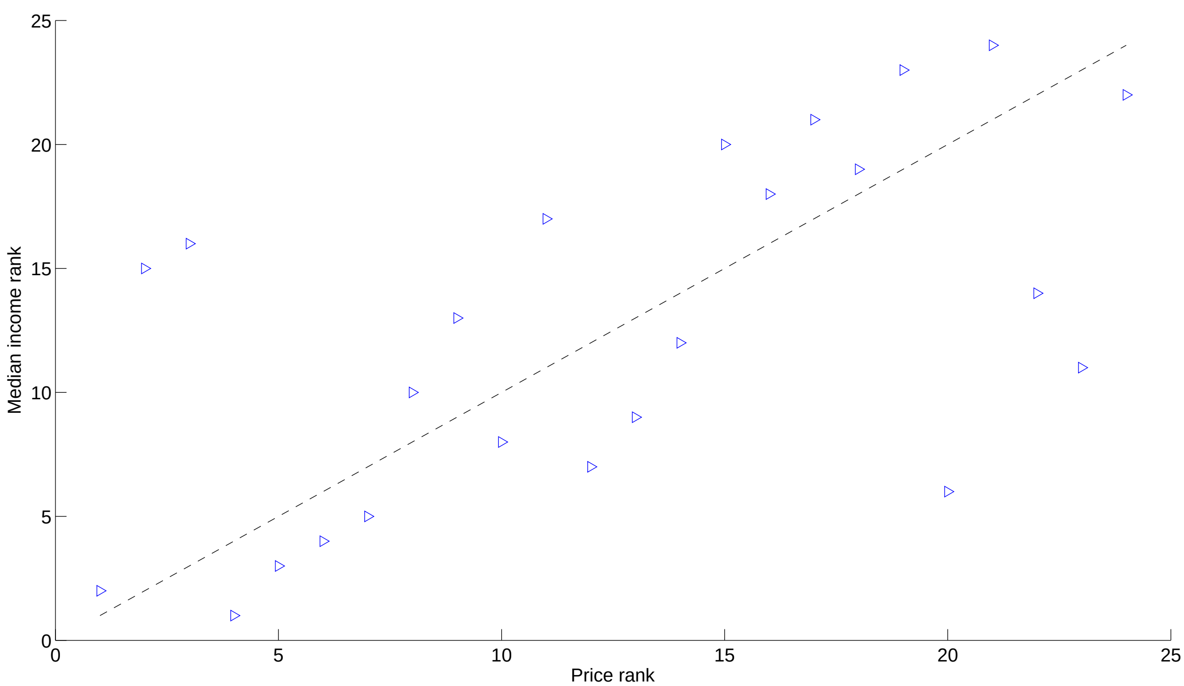

Figure C.12 Log Income - Model I *Out-of-Sample*

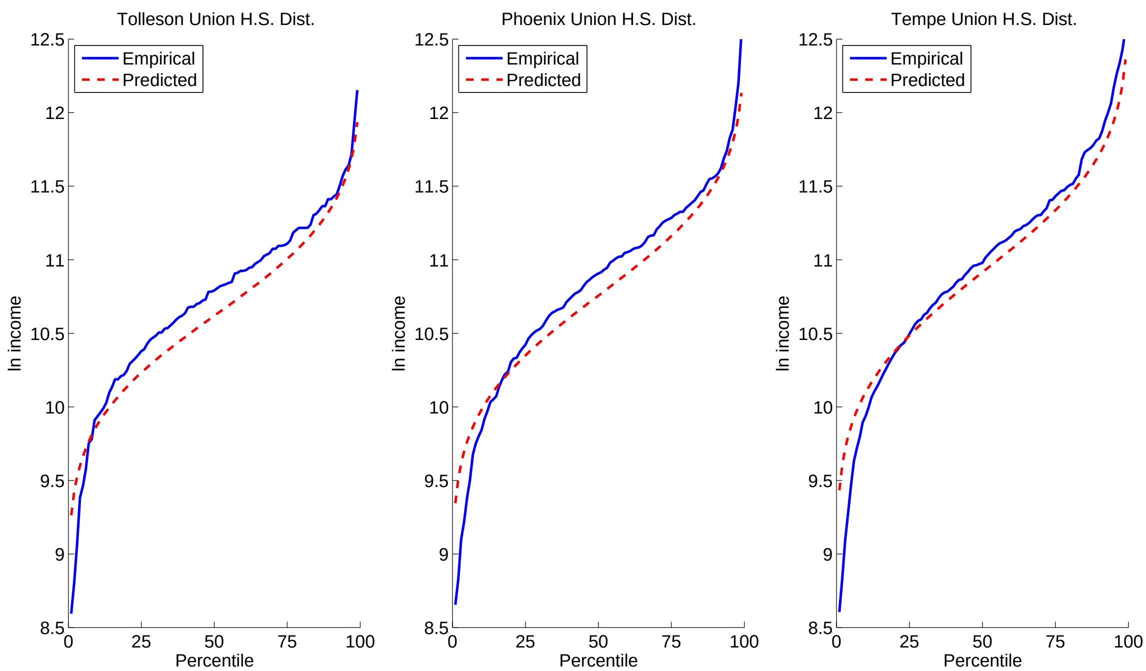

Figure C.13 Log Income - Model II *Out-of-Sample*

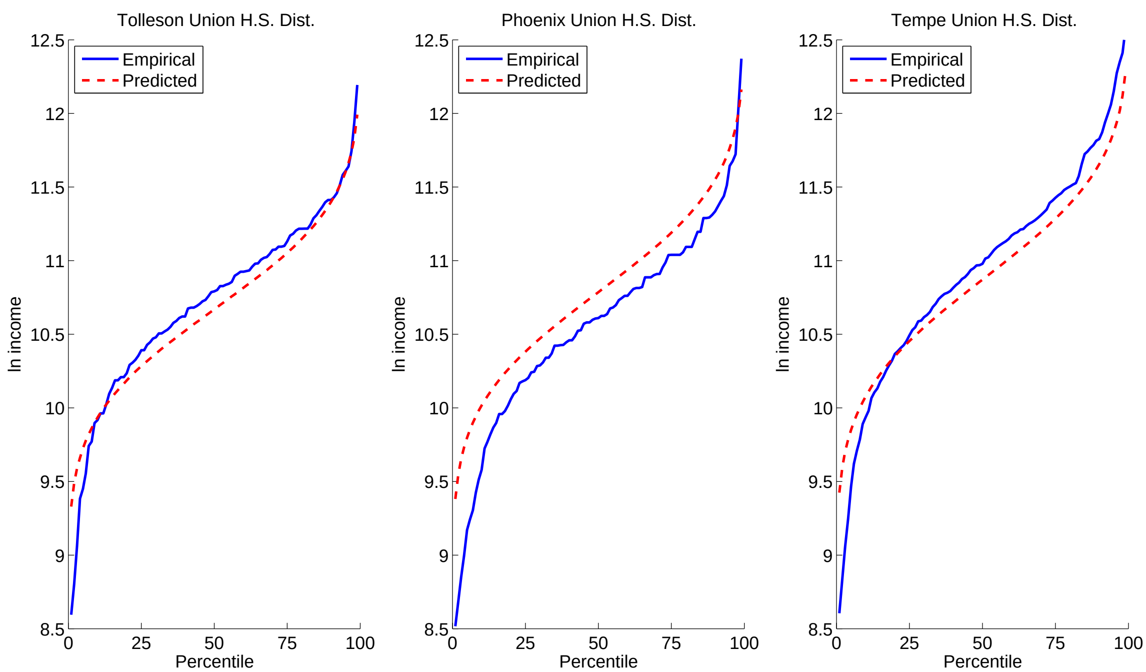

Figure C.14 Log Housing Expenditures - Model I *Out-of-Sample*

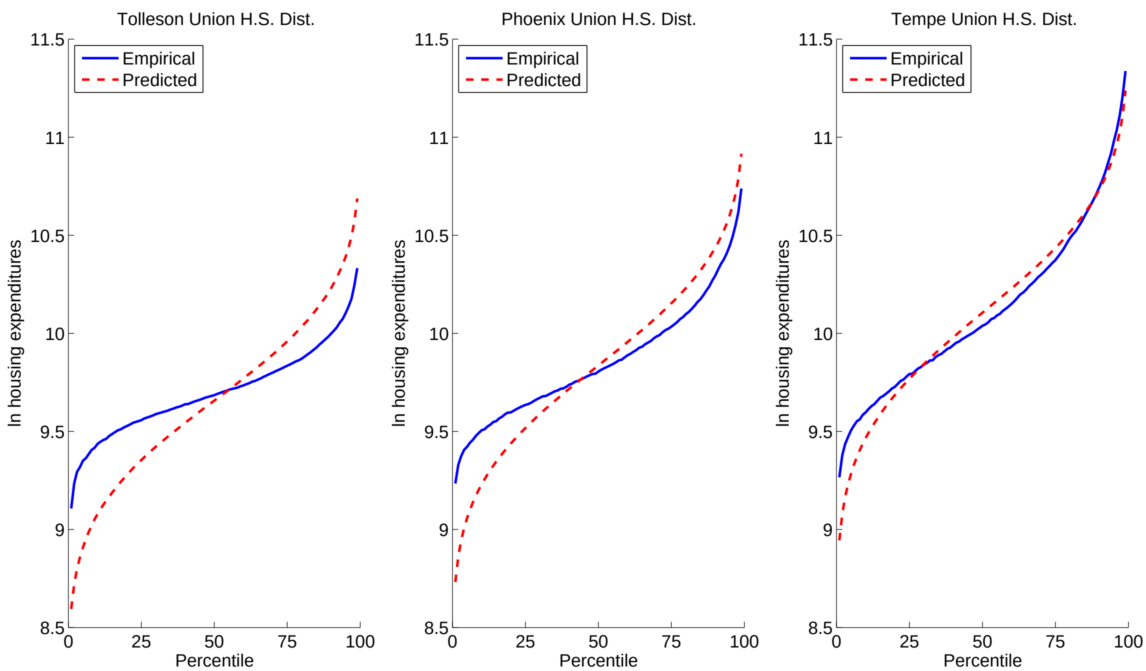

Figure C.15 Log Housing Expenditures - Model II *Out-of-Sample*

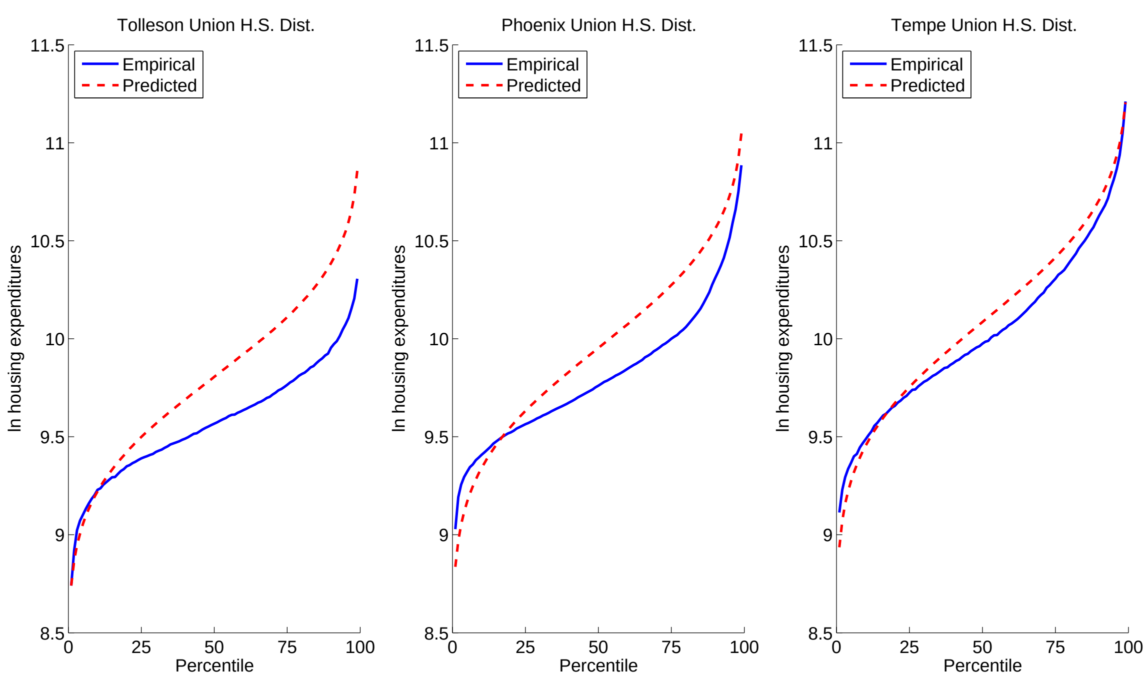

Figure C.16 Prices - Model I *Out-of-Sample*

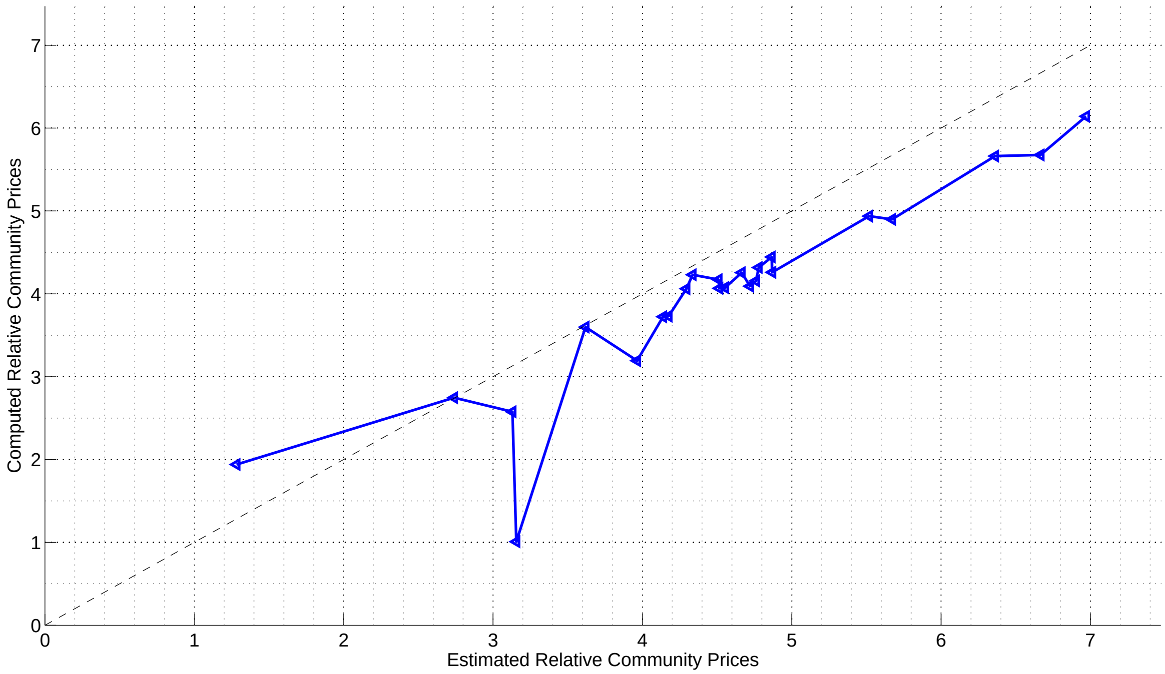

Figure C.17 Prices - Model II *Out-of-Sample*

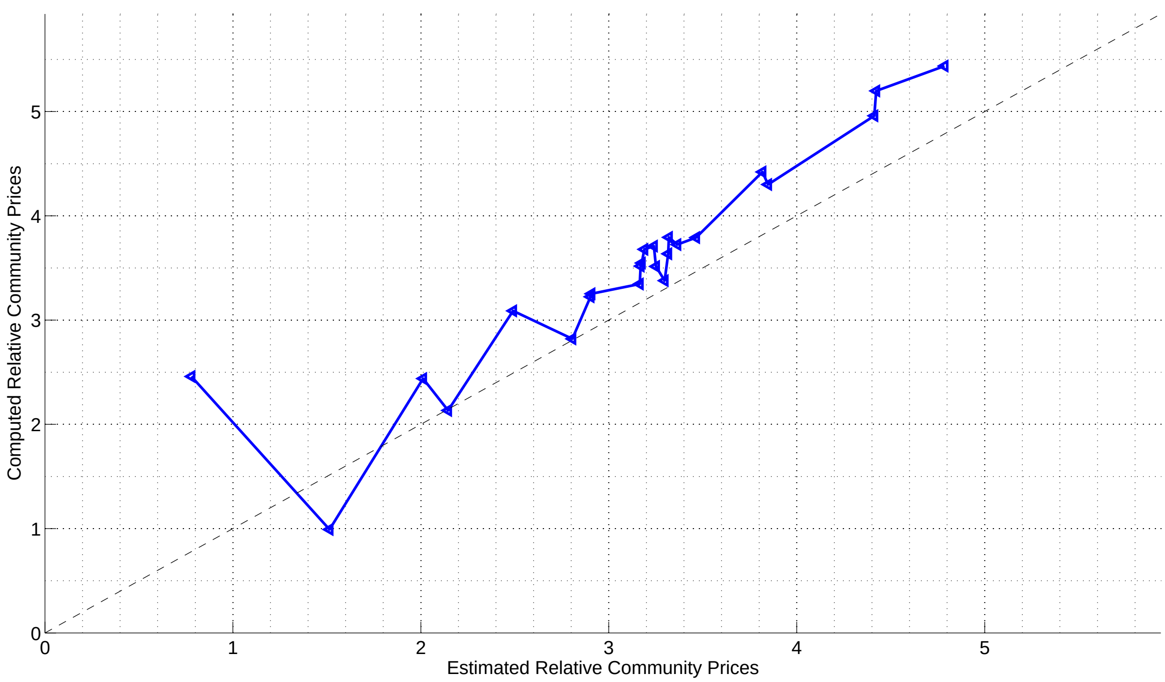

Figure C.18 Population Shares- Model I *Out-of-Sample*

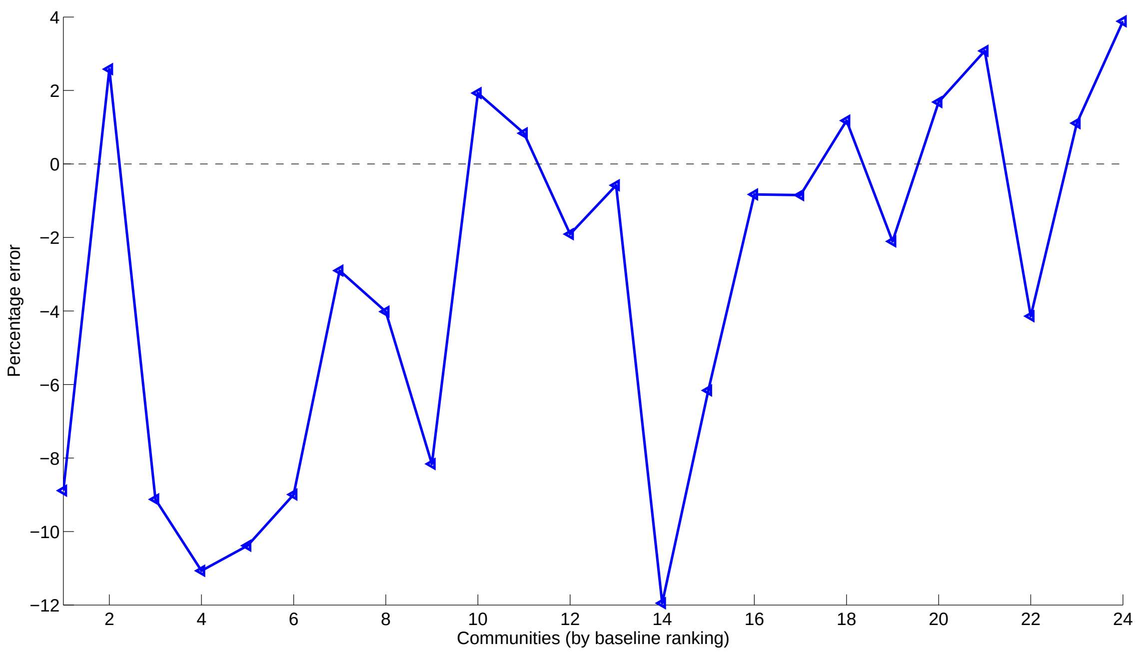

Figure C.19 Population Shares - Model II *Out-of-Sample*

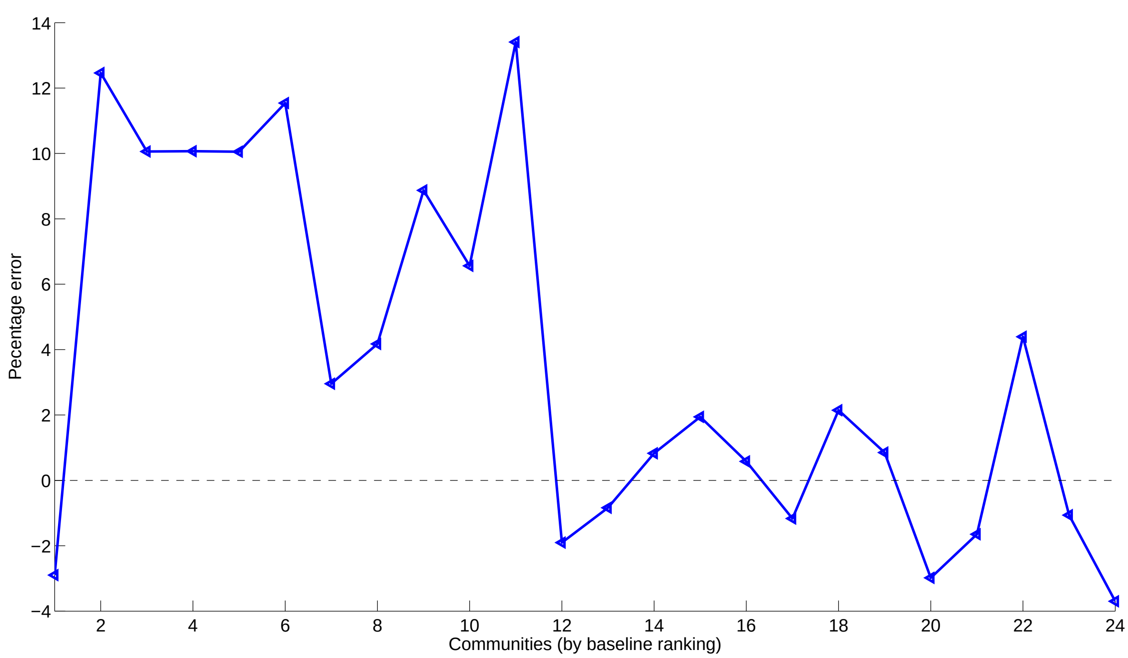

Figure C.20 Price/Public Good Relationship - Model III

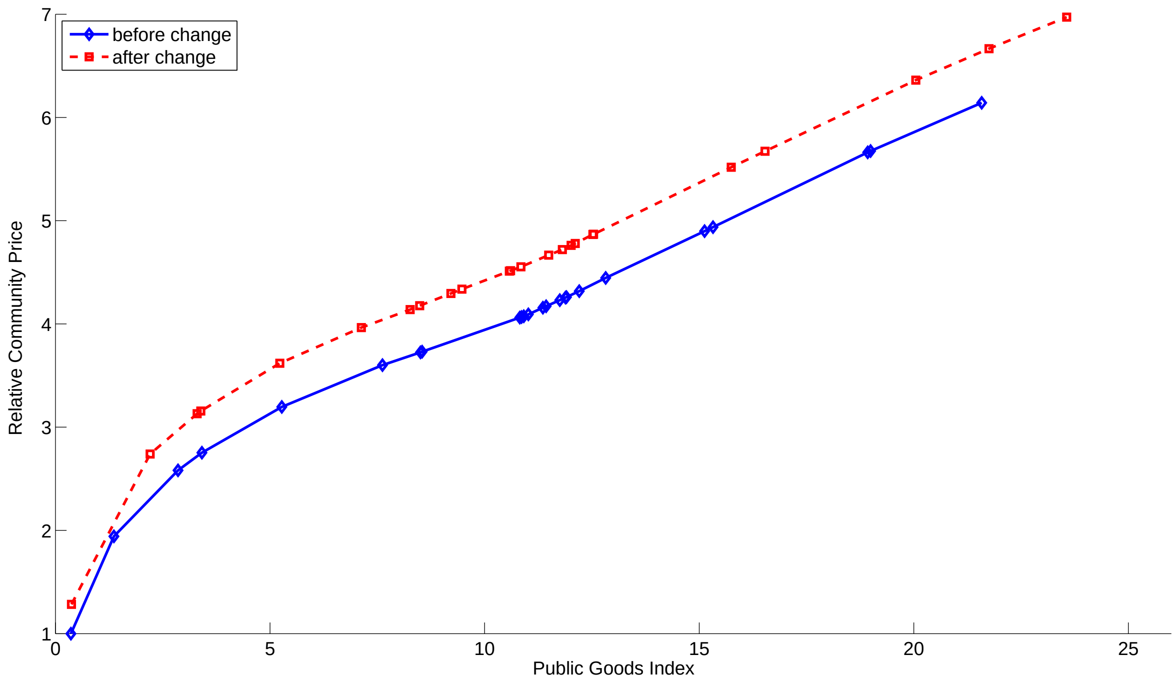

Figure C.21 Mean Willingness-To-Pay for Cleanup of Luke AFB (*Partial Equilibrium*) - Model I

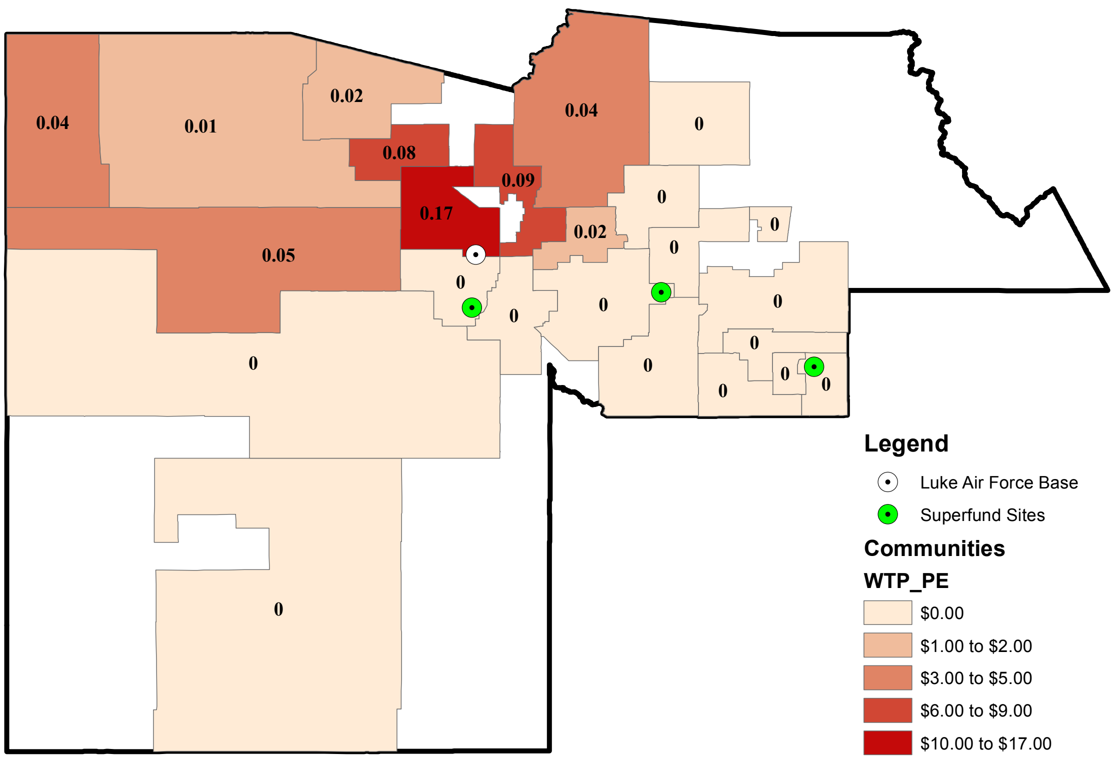

Figure C.22 Mean Willingness-To-Pay for Cleanup of Luke AFB (*General Equilibrium*) - Model I

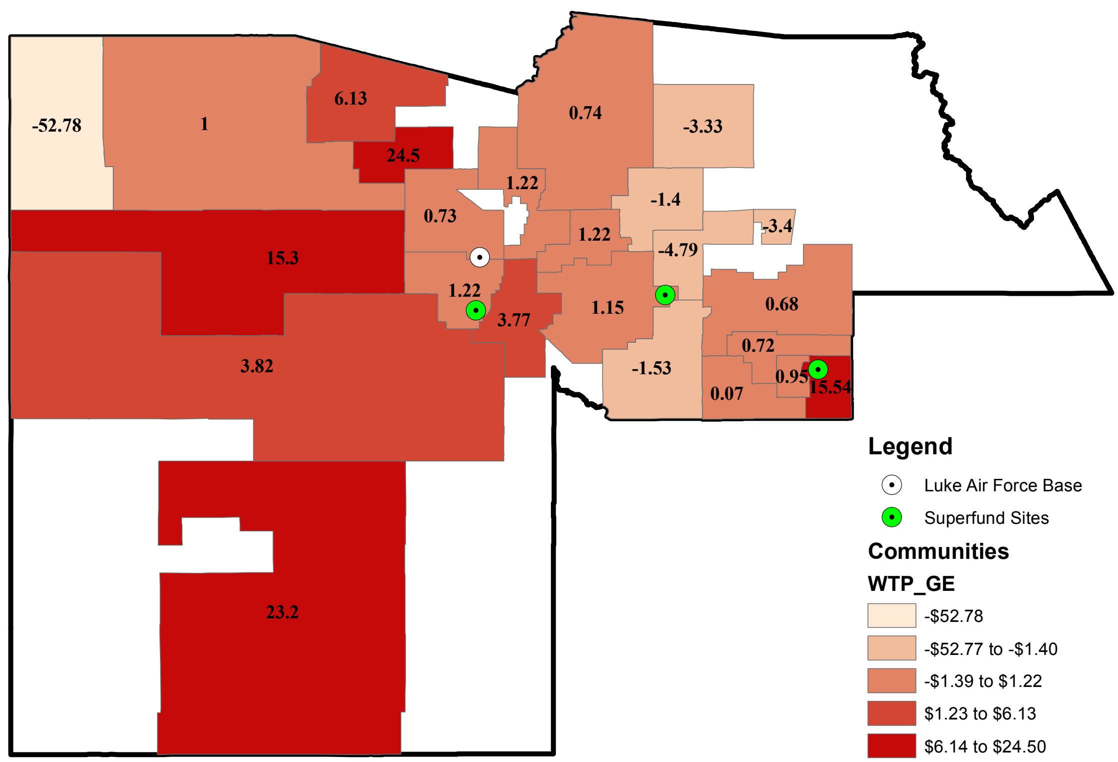

Figure C.23 Mean Willingness-To-Pay for Discovery of Superfund Site (*General Equilibrium*) - Model II