

6-2008

Search and active learning with correlated information: Empirical evidence from mid-Atlantic clam fishermen

Philippe Marcoul
Iowa State University

Quinn Weninger
Iowa State University, weninger@iastate.edu

Follow this and additional works at: http://lib.dr.iastate.edu/econ_las_pubs

 Part of the [Agricultural and Resource Economics Commons](#), [Agricultural Economics Commons](#), and the [Econometrics Commons](#)

The complete bibliographic information for this item can be found at http://lib.dr.iastate.edu/econ_las_pubs/474. For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

This Article is brought to you for free and open access by the Economics at Iowa State University Digital Repository. It has been accepted for inclusion in Economics Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Search and active learning with correlated information: Empirical evidence from Mid-Atlantic clam fishermen

PHILIPPE MARCOUL AND QUINN WENINGER*

March, 2007

Abstract

This paper examines search with active learning and correlated information. We first develop a simple model to show how correlation affects the decision to acquire information. A unique data set on fishing site choice by mid-Atlantic clam fishermen is used to test the model predictions. Results find that clam fishermen search new sites when the catch at familiar sites declines, i.e., when the opportunity cost of gathering information is low, but *also* when catch at familiar sites is on the rise. Search following a catch decline occurs at spatially distant sites whereas search following a catch increase occurs at nearby sites. Correlated learning is crucial for explaining the site choice patterns in our data. These results provide new insights that may extend to a variety of economic search problems where correlated learning is important.

JEL Classification: D8, C4, Q2

*Department of Economics, Iowa State University, email: marcoul@iastate.edu, weningen@iastate.edu. We thank Jacques Crémer, Jérôme Foncel, Corinne Langinier, Taps Maiti, Pau Olivella, Kurt Schmier, Matthew Turner, and seminar participants at the Canadian Resource and Environmental Economics Workshop, Heartland Environmental Workshop, Universities of Québec at Montréal and Wyoming for providing helpful comments. The usual disclaimer applies. Weningen is grateful for financial support provided by the Environmental Protection Agency through the Center for Agricultural and Rural Development.

1 Introduction

Economic problems of search and information acquisition are copious, and a large theoretical literature has studied search under a variety of circumstances. Empirical analysis of search behavior is far less prevalent, particularly in the field. One major obstacle in empirical work is that data on repeated search events intermixed with information acquisition are required to investigate the effects of learning on choice under uncertainty. In addition, the data must include the information that is acquired by the decision maker. For example, an unemployed worker may learn about job prospects through a multitude of sources. Documenting the source and type of information, and the date that information arrives is often not possible. Not surprisingly, evidence of active learning in the field is rare, and questions as to how individuals search and learn in practice remain unanswered.¹

This paper exploits a unique data set on fishing site choices by commercial surf clam fishermen to analyze search and information acquisition in the field, or in our case, on the sea. Clam fishermen make repeated 1-2 day trips to sea to dredge the sea bottom for clams. Skippers cannot know the true clam stock abundance at a selected fishing site until dredging begins and the catch is realized. Fishing a site provides an immediate payoff plus information about the site's true stock abundance. Selecting a site for dredging also implies an opportunity cost equal to the foregone payoff that could be earned at some other possibly more productive site. The problem facing the clam fishermen in our data is similar in many respects to a multiple-armed bandit problem and thus provides a rich setting to examine search and information acquisition.

Clams are sessile creatures and information about true stock abundance at competing sites is valuable; once a productive site is located, a skipper may return repeatedly to *mine* the site's clams. Eventually the vessel's fishing activity will dissipate the local stock abundance and the skipper must again search for a new productive site. As in a bandit problem,

¹Meyer and Shi (1995) and Banks, Olson and Porter (1997) study multi-armed bandit play in the lab. The authors find that decisions by experimental subjects generally follow patterns predicted by models of rational search and learning, although divergence between actual and optimal (rational) play is also found (e.g., Meyer and Shi, 1995).

clam fishermen face a tradeoff between exploration, i.e., dredging at new unfamiliar sites to learn about their true stock abundance, and exploitation, dredging at more familiar sites which are expected to yield a high, but possibly not the highest, payoff. This tension between searching (exploration) and mining (exploitation) inherent in the site choice problem provides the setting for our empirical investigation.

Our data include the catch records for individual vessels on multiple fishing trips; the data spans 15 years with over 1,500 trips for some vessels. Thus we have the requisite longitudinal data on repeated site choice intermixed with information acquisition, i.e., site-specific catch *signals* that skippers use to learn about clam abundance at fished sites. Because knowledge about stock abundance is obtained almost exclusively through the vessel skipper's catch experiences, we observe the information used in the search problem.

There are two dimensions by which the site choice problem facing surf clam skippers in our data differs from a standard bandit problem. First, the standard bandit problem assumes that payoff distributions at competing arms are stationary. We know that payoffs at competing fishing sites change over time as abundance is reduced by dredging and, clams grow. Second, the standard bandit problem assumes that payoff distributions at competing arms are independent. The natural growth characteristics of surf clams imply that stock abundance is spatially correlated. The first part of the paper develops a simple model of search and active learning with correlated information. We show that non-stationarity and correlated learning have important implications for search. The second part of the paper tests the predictions of the model. It turns out that non-stationarity and correlated abundance are important for interpreting search patterns observed in the data.

Our analysis reveals that factors that reduce the cost of acquiring information and factors that enhance the value of information increase the likelihood of a search event, which we define as a decision to dredge for clams at a previously unfished site. Results indicate that search at new sites is more likely when the catch performance at recently-fished sites is on the decline. In the mind of the skipper the cost of search is the foregone payoff from dredging at a new site rather than returning to a familiar site. When the catch performance is on the decline at the familiar site, the cost of gathering information is low making exploration

more attractive. We also find that regulatory changes that increase the value of information increase the likelihood of search. These results are predicted by standard models of search and learning and are not unexpected.

A more surprising result is that we find that exploration at new sites also increases when recent catch performance at familiar sites is on the rise. This result would appear on the surface to contradict optimal search behavior. If the catch rate at a familiar site is on the rise the skipper should perceive a high cost of searching a new site; why switch to a new site when a familiar site is performing well? Our model of search with correlated information explains this search pattern. Because clam stocks are spatially correlated, skippers learn about true abundance at sites that are fished and also at sites believed to be correlated. With correlation an increase in the catch performance at a familiar site raises expectations about abundance at nearby (correlated) sites. A search event that follows a series of above average catch signals may not imply a high search cost at all because with correlation, beliefs about abundance move together. A visit to the nearby unfamiliar site yields valuable information, i.e., the skipper could learn that abundance and thus payoffs are even higher at the nearby site.

Further analysis confirms that if a search event follows a sequence of highly productive trips, it tends to take place at sites that are close in proximity. On the contrary, search events that follow a sequence of low-productivity trips tend to take place at distant or spatially uncorrelated sites. This spatial search pattern is predicted by our model.

We find that search patterns in our data are consistent with a model of rational search with correlated information. Evidence suggests that clam fishermen incur costs in the form of foregone payoffs at familiar sites to explore stock abundance at new sites. We reject naive updating rules in favor of belief updating that geometrically weights past and current information. We also find evidence of risk aversion. Overall, these results suggest that clam skippers actively learn, i.e., they direct their search effort toward sites that yield high value, which is defined as the immediate payoff from fishing plus the value of information.

These results add to the empirical literature on learning. Learning models are common in the marketing literature (Ackerberg, 2001; Erdem and Keane, 1996; see Chintagunta et

al., 2006, for a recent review), and have been used recently to investigate consumer demand for products with uncertain quality characteristics (e.g., Ching, 2005; Crawford and Shum, 2005; Erdem et al., 2004). Much of this literature assumes that learning is independent, that is, that consumer experience with a particular good does not provide information about the quality attributes of other goods in the market. An exception is Erdem (1998) who presents a model with correlated learning across products marketed under a common brand name. An empirical test of the model finds that consumers perceive that the quality of common-brand products is correlated, and correlation affects choices. In Erdem’s model consumers learn passively, i.e., the value of information for guiding future consumption choices is not incorporated into the purchase decision. We extend this literature to consider rational information acquisition in the presence of correlated information. These features may characterize a variety of real world search problems, and our results can shed new light on an important area of research.

The paper is organized as follows. Section 2 presents a two-period model of dynamic site choice. The model demonstrate the key forces that guide the decision to gather information when beliefs are (positively) correlated. The model provides guidance for analysis and interpretation of our data. Section 3 provides background information for the surf clam fishery and presents descriptive statistics. Empirical analysis and results are presented in Section 4. Section 5 summarizes our findings and provides concluding remarks.

2 A simple model

The empirical investigation will focus on the trade-off between exploration and exploitation, which in the context of a fishing site choice can be framed as a decision to explore a new site or return to a familiar site. We develop the simplest possible model to characterize this trade-off.²

²Non-stationarity and correlated abundance (payoffs) rules out use of a reservation price solution approach (Gittins, 1979). The multiple-site, multiple-period problem is intractable. With K sites, the value function and site choice policy function depend on $\frac{1}{2}K(K + 3)$ continuous state variables, i.e., the site-specific beliefs about mean payoffs, plus the variance and covariance terms. There are in excess of 200 sites

Consider a single fisherman who fishes at three separate fishing sites, denoted a , b and c , during two sequential periods. Sites are selected at dates t_0 and t_1 . Fishing takes place between t_0 and t_1 and following date t_1 . Only one site can be fished in each period. We assume fishing at some site is always preferred to not fishing at all. A fishing strategy is a sequence of chosen sites; under strategy (a, b) , site a is fished in the first period and b is fished in the second. In our simple model, there are ex ante nine possible strategies.

True stock abundance is represented by positive real numbers, denoted α for site a , β for site b , and γ for site c . To simplify notation, we assume that payoffs from fishing are proportional to stock abundance and, henceforth, we refer to payoffs and stock abundance synonymously.

We assume the fisherman has prior beliefs about abundance at each site but is never certain about true abundance or true payoffs. To represent this uncertainty, assume the fisherman believes abundance at each site follows a normal distribution. The normal distribution offers a computationally convenient framework to represent and analyze belief formation. While it is clear that stock abundance cannot be negative, the net returns on a fishing trip may be, and thus, normally distributed *payoffs* are not unreasonable. We will proceed under the interpretation that model variables represent stock abundance or payoffs, noting that the results that follow are robust to the case where the belief distribution has support on the negative orthant (multiplicative stock depletion would require modification in the case of negative payoffs). At date t_0 , the beliefs about abundance are normally distributed following

$$(1) \quad \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} \sim N \left[\begin{pmatrix} \alpha_0 \\ \beta_0 \\ \gamma_0 \end{pmatrix}, \begin{pmatrix} \sigma_\alpha^2 & \sigma_{\alpha\beta} & 0 \\ \sigma_{\alpha\beta} & \sigma_\beta^2 & 0 \\ 0 & 0 & \sigma_\gamma^2 \end{pmatrix} \right],$$

where α_0 , β_0 and γ_0 denote means, and σ_α^2 , σ_β^2 and σ_γ^2 denote variances at sites a , b and c ,

in our data. If we assume only the top 10% of these sites are viable choices, the number of state variables would be 230. This curse of dimensionality complicates dynamic programming solutions to the multiple-site multiple-period problem and rules out the use of numerical methods (Gabaix and Laibson, 2000).

respectively. The variance covariance matrix is assumed positive definite. We use $\rho \equiv \frac{\sigma_{\alpha\beta}}{\sigma_\alpha\sigma_\beta}$ to denote the correlation coefficient.

The realized catch provides information about the site's true stock abundance. We assume that fishing at site a , b or c in period 1 yields a catch *signal* of the form

<i>Site</i>	<i>Signal</i>
a	$S_a = \alpha + \varepsilon_a$
b	$S_b = \beta + \varepsilon_b$
c	$S_c = \gamma + \varepsilon_c$.

Each signal contains a noise component, respectively ε_a , ε_b , and ε_c at sites a , b and c , which is assumed independent and normally distributed with zero mean and known variance, denoted σ_s^2 .³ Uncertainty in the catch signal reflects randomness with which the fishing gear intercepts the stock. It is reasonable to expect that σ_s^2 is small relative to the perceived uncertainty in true stock abundance, and we assume $\sigma_s < \sigma_j$, $j = \alpha, \beta$, and γ . Notice that realized signals are also the payoffs from fishing at a site.

Finally, we allow harvest activity to deplete the stock. Let $\delta \in [0, 1]$ denote the fraction of the expected abundance that is available at a site which has been fished in period 1. As δ tends to zero, first period harvesting fully depletes the stock at the site; $\delta = 1$ indicates no stock depletion. There is no natural stock growth in the model and no discounting but each of these elements could be introduced.⁴

When $\delta = 1$ and $\rho = 0$, the fisherman faces a standard two-period, three-armed bandit problem. When $\delta \in [0, 1)$ the site choice problem is signal-dependent since the period 2 choice set depends on the signal that is obtained in the first fishing period (Sulganik and Zilcha, 1997). Note that if $\delta = 0$ and $\rho = 0$, there is no useful learning in the model, and information plays no role in guiding future site choice.

Belief updating

³Site-specific signal variance is easily incorporated into the model but adds few additional insights.

⁴ δ is a parameter in the decision problem. Formally incorporating δ into the belief structure of the fisherman would further complicate the analysis leaving the main results qualitatively unchanged.

At date t_1 , the fisherman has acquired a signal that will be used to update beliefs about true stock abundance. We assume the fisherman updates beliefs following Bayes rule. Suppose correlated site a is fished in period 1. Conditional on the signal S_a , updated beliefs at a follow a normal distribution (De Groot, 1970):

$$(2) \quad (\alpha|S_a) \sim N \left(\delta \frac{S_a \sigma_\alpha^2 + \alpha_0 \sigma_s^2}{\sigma_\alpha^2 + \sigma_s^2}, \delta^2 \frac{\sigma_\alpha^2 \sigma_s^2}{\sigma_\alpha^2 + \sigma_s^2} \right).$$

Updated mean abundance is a convex combination of the prior belief and the signal S_a . Ignoring stock depletion for a moment ($\delta = 1$), good signals $S_a > \alpha_0$ increase expected abundance at site a . If the catch signal is noisy, i.e., σ_s^2 is large, the prior mean α_0 receives more weight, whereas with low signal noise more weight is given to the new information. Of course, the posterior mean does not change if $S_a = \alpha_0$. All signals reduce the posterior variance by the factor $\sigma_s^2/(\sigma_\alpha^2 + \sigma_s^2)$ which is less than unity for $\sigma_\alpha^2, \sigma_s^2 > 0$. With stock depletion, $\delta < 1$, the mean and variance are reduced multiplicatively by δ and δ^2 , respectively.

Consistent with our empirical application we focus on non-negative correlation $\rho \in [0, 1)$ (negative correlation is considered in an extended appendix). With correlated beliefs, catch signals are used to update beliefs at unfished sites. Conditional on S_a , updated beliefs at b are given as

$$(3) \quad (\beta|S_a) \sim N \left(\beta_0 + \rho \frac{\sigma_\alpha \sigma_\beta}{(\sigma_\alpha^2 + \sigma_s^2)} (S_a - \alpha_0), \sigma_\beta^2 \left(1 - \rho^2 \frac{\sigma_\alpha^2}{(\sigma_\alpha^2 + \sigma_s^2)} \right) \right).$$

With $\rho > 0$ a favorable signal, $S_a > \alpha_0$, leads to an upward adjustment in mean abundance at b , relative to the prior mean β_0 . The adjustment is proportional to $S_a - \alpha_0$ and increases with ρ . As above, with larger signal noise σ_s^2 , less weight is given to the catch signal.⁵ Expression (3) does not depend on δ since the existence of S_a implies site b has not been fished. It is easy to see that with $\rho = 0$ no new information about abundance b is obtained and the posterior mean is unchanged. Updated beliefs conditional on a first period signal from site b are symmetric and are presented in an appendix.

⁵The posterior covariance conditional on obtaining signal S_a and S_b is, respectively, $\rho \sigma_\alpha \sigma_\beta \sigma_s^2/(\sigma_\alpha^2 + \sigma_s^2)$ and $\rho \sigma_\alpha \sigma_\beta \sigma_s^2/(\sigma_\beta^2 + \sigma_s^2)$.

Fishing at uncorrelated site c in period one yields S_c , but no information about abundance at a and b and consequently only site c beliefs are updated:

$$(4) \quad (\gamma|S_c) \sim N \left(\delta \frac{S_c \sigma_\gamma^2 + \gamma_0 \sigma_s^2}{\sigma_\gamma^2 + \sigma_s^2}, \delta^2 \frac{\sigma_\gamma^2 \sigma_s^2}{\sigma_\gamma^2 + \sigma_s^2} \right).$$

Assume the fisherman is risk neutral. The objective is to maximize the two-period payoffs given prior beliefs about abundance. To simplify the intuition, we will initially remove the uncorrelated site c from the choice set. Site c is reintroduced below.

Consider the decision of where to fish in period 2. Suppose site a was fished in the first period. The fisherman returns to a only if updated beliefs satisfy $E[\alpha|S_a] \geq E[\beta|S_a]$, where E is the expectation operator. From (2) and (3), a return trip to a is warranted if

$$(5) \quad \delta \frac{S_a \sigma_\alpha^2 + \alpha_0 \sigma_s^2}{\sigma_\alpha^2 + \sigma_s^2} \geq \beta_0 + \rho \frac{\sigma_\alpha \sigma_\beta}{(\sigma_\alpha^2 + \sigma_s^2)} (S_a - \alpha_0).$$

The fisherman will switch to b in period 2 if $E[\beta|S_a] > E[\alpha|S_a]$. We define a threshold signal, denoted \underline{S}_a , at which he is indifferent between site a and b ; \underline{S}_a solves $E[\alpha|\underline{S}_a] = E[\beta|\underline{S}_a]$. If b is fished in period 1, the fisherman will return in period 2 only if updated beliefs satisfy $E[\beta|S_b] \geq E[\alpha|S_b]$. There will exist a threshold \underline{S}_b that solves $E[\beta|\underline{S}_b] = E[\alpha|\underline{S}_b]$ (solutions for \underline{S}_a and \underline{S}_b appear in an Appendix).

With only two sites, a and b , the optimal policy at date t_1 can be expressed as a function of the period 1 signal: conditional on S_a (i.e., having fished at a in period 1), fish at a in period 2 if $S_a \geq \underline{S}_a$, otherwise switch to b ; and, conditional on S_b (i.e., having fished at b in period 1), fish at b in period 2 if $S_b \geq \underline{S}_b$, otherwise switch to a .

Now consider the date t_0 choice of the first period site. Our assumptions about beliefs imply that the date t_0 marginal probability distribution of signals is normal (see Zellner, 1971), e.g., $S_a \sim N(\alpha_0, \sigma_\alpha^2 + \sigma_s^2)$. With the date t_1 policy summarized above we calculate the date t_0 probability distribution of fishing at each site in period 2.

Suppose site a is fished first. Define the standardized normal random variable $\lambda_a \equiv (S_a - \alpha_0)/\sqrt{\sigma_\alpha^2 + \sigma_s^2}$. The t_0 probability of making a return trip to a in period 2 is $\Pr(S_a > \underline{S}_a) = 1 - \Phi(\underline{\lambda}_a)$, where $\underline{\lambda}_a \equiv (\underline{S}_a - \alpha_0)/\sqrt{\sigma_\alpha^2 + \sigma_s^2}$ and Φ is the standard normal cumulative

distribution function. The t_0 probability of switching to site b in period 2 is $\Phi(\underline{\lambda}_a)$. Next define the expectation of truncated signals,

$$\begin{aligned} S_a^+ &= E[S_a | S_a > \underline{S}_a] = \alpha_0 + \sqrt{(\sigma_\alpha^2 + \sigma_s^2)} \frac{\phi(\underline{\lambda}_a)}{1 - \Phi(\underline{\lambda}_a)} \\ S_a^- &= E[S_a | S_a < \underline{S}_a] = \alpha_0 - \sqrt{(\sigma_\alpha^2 + \sigma_s^2)} \frac{\phi(\underline{\lambda}_a)}{\Phi(\underline{\lambda}_a)}, \end{aligned}$$

where ϕ is the standard normal probability distribution. The date t_0 expected payoff from fishing at site a in period one is

$$(6) \quad V(a) = \alpha_0 + [1 - \Phi(\underline{\lambda}_a)]E[\alpha | S_a > \underline{S}_a] + \Phi(\underline{\lambda}_a) E[\beta | S_a < \underline{S}_a],$$

where $E[\alpha | S_a > \underline{S}_a]$ is obtained from (2) using expected signal S_a^+ , and $E[\beta | S_a < \underline{S}_a]$ is obtained from (3) with expected signal S_a^- . $V(a)$ is easily interpreted: the fisherman expects payoff α_0 in period 1, and optimally responds to the signal S_a when selecting the second period site. With probability $[1 - \Phi(\underline{\lambda}_a)]$ he expects to return to a and obtain expected payoff $E[\alpha | S_a > \underline{S}_a]$. With probability $\Phi(\underline{\lambda}_a)$ site a yields an *unfavorable* signal in which case he switches to b and collects $E[\beta | S_a < \underline{S}_a]$ in period 2. Similar steps are used to calculate the date t_0 expected payoff from fishing first at site b :

$$(7) \quad V(b) = \beta_0 + [1 - \Phi(\underline{\lambda}_b)]E[\beta | S_b > \underline{S}_b] + \Phi(\underline{\lambda}_b) E[\alpha | S_b < \underline{S}_b].$$

The optimal strategy at date t_0 is to fish at a first if $V(a) > V(b)$, otherwise fish at b first. The fisherman is indifferent as to which site is fished first if $V(a) = V(b)$.

We next investigate the incentive to acquire information. Exploration at uncertain sites provides information that guides future site choice, and an optimal search strategy may involve an investment in information, i.e., visiting a site with a lower expected immediate payoff but with high information value. That is, it may be optimal to go through a “payoff valley” to potentially reach a “payoff mountain” later on (Adam, 2001, p. 264). The next section examines the role of correlation in the decision to *invest* in information.

2.1 Information acquisition

To further illustrate the trade-off between exploration and exploitation, for the remainder of the paper site a will play the role of a *familiar* site in the sense that the fisherman is assumed to be most certain about true abundance at site a , i.e., $\sigma_\alpha^2 < \sigma_\beta^2$. Suppose also that date t_0 beliefs about abundance at sites a and b satisfy $\alpha_0 > \beta_0$ and $\delta = 1$ (optimal site choice with $\delta < 1$ is examined below). With these beliefs the decision to fish at b represents an investment in information since fishing at b implies a first-period foregone expected payoff $\alpha_0 - \beta_0 > 0$. The following proposition characterizes the optimal fishing strategy as a function of the uncertainty at unfamiliar site b . This result is standard and has been derived in quite general settings with $\rho = 0$ (e.g., Adam, 2001).

Proposition 1 (Investment in information) *When the beliefs of a risk neutral fisherman satisfy $\delta = 1, \alpha_0 > \beta_0$, there exists a unique threshold value σ_β^2 , denoted $\bar{\sigma}_\beta^2$, such that for any $\sigma_\beta^2 > \bar{\sigma}_\beta^2$, the fisherman prefers to fish at site b in period 1. When $\sigma_\beta^2 < \bar{\sigma}_\beta^2$, site a is preferred in period 1.*

Proof The proof is in an appendix. ■

When σ_β^2 is sufficiently high, it pays to explore site b to acquire information. This result holds true in the presence of correlation across sites, $\rho \neq 0$. When σ_β^2 is large, there is a chance of a very high and low catch at b . The option to switch sites in the second period truncates the period 2 expected payoff from below. Information acquired while fishing the unfamiliar site can raise the date t_0 expected reward, making a trip to the uncertain site, or an investment in information, worthwhile. As uncertainty at b increases the fishermen is willing to incur a higher first period cost to learn about site b 's true reward. On the contrary, as $\alpha_0 \rightarrow \beta_0$ from above, the threshold value $\bar{\sigma}_\beta^2 \rightarrow \sigma_\alpha^2$ (see Appendix for proof); when the cost of gathering information is small a fisherman is more willing to visit an unfamiliar site to learn about its true abundance.

Correlation across sites influences the incentive to gather information. The larger is the covariance term $\sigma_{\alpha\beta}$, the larger must be the threshold $\bar{\sigma}_\beta^2$ (this claim is easily verified using proposition 1 and proposition 2 below). Suppose the fisherman visits the unfamiliar site b

first. The period 2 expected payoff conditional on the signal received is

$$(8) \quad \max (E[\alpha|S_b], E[\beta|S_b]) = \max \left(\alpha_0 + \rho \frac{\sigma_\alpha \sigma_\beta}{(\sigma_\beta^2 + \sigma_s^2)} (S_b - \beta_0), \frac{S_b \sigma_\beta^2 + \beta_0 \sigma_s^2}{\sigma_\beta^2 + \sigma_s^2} \right).$$

With $\rho = 0$, the signal S_b provides no off-site information and expected abundance at a , the first right-hand term in (8), remains α_0 . With $\rho > 0$ above (below) average signals raise (lower) expectations about abundance at both sites. Hence, period 2 updated beliefs move together and guidance provided by the signal S_b as to where to fish in period 2 is less clear cut. Moreover, when abundance is positively correlated the fisherman believes, at date t_0 , that with some probability he will be *skunked* in period 2, i.e., abundance and catch will be low at both sites. The value of the option to switch sites in response to an unfavorable first period signal is smaller when $\rho > 0$, and as a result the fisherman will be less willing to incur a period 1 cost to gather information. This intuition is summarized in the following proposition.

Proposition 2 *Assume that an investment in information is profitable, $V(b) > V(a)$. The relative payoffs from the investment, measured as $\Delta = V(b) - V(a)$, decrease as the covariance of stock abundance $\sigma_{\alpha\beta}$ increases.*

Proof The proof is in the appendix. ■

A rational fisherman realizes that information gathered while fishing has value, and includes this value in the first period site choice. In contrast a myopic searcher will select site a in period 1 if $\alpha_0 > \beta_0$. Period 2 beliefs are updated conditional on the signal S_a received in the first period yielding expected payoff

$$V(a) = \alpha_0 + [1 - \Phi(\underline{\lambda}_a)] E[\alpha|S_a > \underline{S}_a] + \Phi(\underline{\lambda}_a) E[\beta|S_a < \underline{S}_a].$$

By Proposition 1 however, $V(a) < V(b)$ for all $\sigma_\beta^2 > \bar{\sigma}_\beta^2$, which implies different strategies for rational and myopic searchers. When $\sigma_\alpha^2 \approx \sigma_\beta^2$, both a myopic fisherman and an active learner will prefer the higher mean site a . However, from proposition 1 when $\sigma_\beta^2 \gg \sigma_\alpha^2$, a forward looking fisherman will prefer to gather information that can guide future site choice.




Figure 1: Period 2 updated beliefs conditional on period 1 signal, S .

2.1.1 Directed search: correlated versus independent sites

We now reintroduce the uncorrelated site c to the choice set. Adding an independent site allows us to consider the decision of *where* to acquire information and does not qualitatively alter the results obtained so far. Denote the date t_0 expected payoff from fishing at site c first as $V(c)$ which is calculated similarly to $V(a)$ and $V(b)$.

With three sites available the decision of where to acquire information arises. To investigate this choice, we first isolate the value of information obtained at a correlated versus an independent site. Assume that initial beliefs satisfy $\alpha_0 = \beta_0 = \gamma_0$ and $\sigma_\alpha^2 = \sigma_\beta^2 = \sigma_\gamma^2$, with $\sigma_{\alpha\beta} > 0$, and consider the value information from site b versus site c . Figure 1 shows period 2 updated beliefs as a function of the period 1 signal. The vertical axes show the updated mean payoffs; first period signals appear on the horizontal axes. For notational ease and to avoid confusion, signal subscripts are dropped. The assumptions for prior beliefs imply that the period 1 expected payoff and the date t_0 signal distribution is the same across sites. This enables us to characterize $V(\cdot)$ by focussing on period 2 payoffs conditional on the period 1 signal.

The left-hand panel of Figure 1 depicts updated beliefs for the case where independent site c is fished first. The right-hand panel shows updated beliefs when b is fished first (under the assumed initial beliefs the case where a is fished first is symmetric and is not considered).

Assume initially that $\delta = 1$ (no stock depletion) and suppose site c is fished first. The

updated belief about mean abundance is denoted $E[\gamma|S]_{\delta=1}$. A signal $S < \gamma_0$ will cause the fisherman to switch to site b (or a) in period 2, in which case the period 2 expected payoff is β_0 (α_0). Consequently if c is fished first, period 2 expected payoffs are the maximum of $E[\gamma|S]_{\delta=1}$ and β_0 . Similar logic finds that if site b is fished first, period 2 expected payoff is the maximum of $E[\beta|S]_{\delta=1}$ and γ_0 . By construction $E[\gamma|S]_{\delta=1} = E[\beta|S]_{\delta=1}$ and $\gamma_0 = \beta_0$, and because the date t_0 signal distribution is identical at b and c it is easily seen that $V(c) = V(b)$ ($= V(a)$) when $\rho > 0$, and $\delta = 1$.

Although a visit to site b provides *extra information* about abundance at a , the figure shows that this information is not valuable since any signal $S > \beta_0$ increases the updated mean at b more than at a , i.e., $E[\beta|S > \beta_0]_{\delta=1} > E[\alpha|S > \beta_0]$. The optimal policy will never involve a visit to b when $S > \beta_0$. Furthermore, when $S < \beta_0$ the fisherman prefers the independent site c and a return to b is never optimal.

The extra or correlated information can have value when $\delta < 1$. In Figure 1, $E[\gamma|S]_{\delta<1}$ and $E[\beta|S]_{\delta<1}$ denote updated beliefs at site c and b , respectively, when $\delta < 1$. If c is fished first, the fisherman will return to obtain $E[\gamma|S]_{\delta<1}$ when $S \geq \bar{S}_{c,\delta<1}$. Otherwise, the fisherman switches to b and receives β_0 . Compare these payoffs to those obtained when b is fished first. A signal $S_b \geq \bar{S}_{b,\delta<1}$ will prompt a return to b yielding $E[\beta|S_b]_{\delta<1}$. If $S_b < \beta_0$, the fisherman switches to c in period 2 and receives γ_0 . For signals in the interval $[\beta_0, \bar{S}_{b,\delta<1}]$ the fisherman switches to a and receives $E[\alpha|S_b]$ in period 2. Because abundance is believed positively correlated, $E[\alpha|S_b] > \beta_0$ for $S_b \in [\beta_0, \bar{S}_{b,\delta<1}]$. Once again, the period 1 expected payoff, and the date t_0 signal distribution is the same at c and b ; it is easily seen that $V(b) > V(c)$ when $\rho > 0$ and $\delta < 1$. From this we conclude that without depletion correlated information has no value to the fisherman.

Lastly, notice that as $\sigma_{\alpha\beta}$ and thus ρ becomes larger, $E[\alpha|S_b]$ pivots counterclockwise around point m (this rotation will also increase $\bar{S}_{b,\delta<1}$). This suggests that with $\delta < 1$ an increase in $\sigma_{\alpha\beta}$ will increase $V(b)$ relative to $V(c)$.

2.1.2 More on directed search

Let site a continue to play the role of a familiar (low variance) site. We focus attention on date t_0 beliefs that satisfy $\sigma_\alpha^2 < \sigma_\beta^2 \leq \sigma_\gamma^2$, where the relation $\sigma_\beta^2 \leq \sigma_\gamma^2$ follows under correlated abundance at a and b , and under the assumption that a has the lowest variance. Assume again that the fisherman believes a trip to b or c represents an investment in information in the sense that $\alpha_0 > \beta_0$ and $\alpha_0 > \gamma_0$. If a is believed to have the highest abundance and $\rho > 0$ then it is also reasonable that $\beta_0 \geq \gamma_0$. The following proposition identifies beliefs for which an investment in information at the correlated site is preferred to an investment in information at the uncorrelated site.

Proposition 3 (Information at correlated and independent sites) *Assume the beliefs of a risk neutral fisherman satisfy $\alpha_0 > \beta_0 > \gamma_0$. By Proposition 1, there exists a $\bar{\sigma}_\gamma^2$ for which the fisherman is indifferent between visiting site b and site c in the first fishing period. If beliefs also satisfy $\sigma_\alpha^2 < \sigma_\beta^2 \leq \sigma_\gamma^2 < \bar{\sigma}_\gamma^2$, there exists a $\tilde{\beta}_0$ such that*

- (I) *for $\beta_0 < \tilde{\beta}_0$ investment in information, if it occurs, will occur at site c ;*
- (II) *for $\beta_0 \geq \tilde{\beta}_0$, investment in information, if it occurs, will occur at site b .*

The proof of Proposition 3 has a straightforward intuition. First, notice that if $\beta_0 \rightarrow \gamma_0$, exploration, if it occurs, should be at site c . This follows from the fact that site b and c offer similar mean abundance, while fishing at the higher variance site c yields more information (Proposition 1). When $\beta_0 \rightarrow \alpha_0$, site a becomes a dominated choice and the fisherman will prefer to visit either b or c first. By Proposition 1 there will exist a threshold value $\sigma_\gamma^2 = \bar{\sigma}_\gamma^2$ such that a visit to c is not warranted if $\sigma_\gamma^2 < \bar{\sigma}_\gamma^2$, which is assumed in the proposition. The continuity of the functions $V(a)$, $V(b)$, and $V(c)$ in β_0 insures the existence of $\tilde{\beta}_0$, which completes the proof.

It is clear that the decision to explore hinges on the perceived costs and benefits of gathering information. If $\beta_0 = \gamma_0$, the cost of gathering information at b and c is the same and the higher information value at the independent site tips the scale in its favor. Proposition 3 states simply that information acquisition at the correlated site will be preferred only if the cost of gathering the information at the site, $\alpha_0 - \beta_0$, is sufficiently small.

2.1.3 Implications for multiple-site, multiple-period search

Before we turn to the empirical analysis we consider how the insights gained from the simple model can be extended to the case of multiple sites and multiple periods. The discussion is necessarily informal.

Our empirical application exploits the unique feature of clam fishing that once a productive site is located skippers make repeated trips to *mine* its clams (89.82% of the trips are of this type). Notice that a skipper who has visited a site has obtained a catch signal from the site. It is reasonable to suspect that the skipper is more certain about true stock abundance at previously fished sites than at sites that have not been fished. In the mind of a skipper, previously fished sites should represent relatively certain payoff (low variance) sites whereas unfished sites should represent relatively uncertain (high variance) payoff sites. In the context of our model, fished sites correspond to the familiar site a .

In the multiple site case, skippers choose the distance between sites. Because spatial correlation declines with distance, choosing a site that is close to a familiar site is tantamount to choosing a positively correlated site, i.e., site b from the model. Alternatively choosing to fish at a site that is distant from a familiar site, is tantamount to fishing at an independent site, site c in the model.

In the context of the model, a skipper who makes a repeat trip to a familiar site prefers a over b or c . The skipper should believe that abundance at the site being dredged exceeds abundance at other available sites, otherwise some other site would have been selected for dredging. Alternatively, the fact that a skipper visited a familiar rather than a new site suggests that the perceived value of information does not offset the perceived cost of gathering the information. The beliefs of a skipper on repeat trips to a familiar site are likely characterized as:

$$(9) \quad \begin{aligned} \alpha_0 &> \beta_0 \geq \gamma_0 \\ \sigma_\alpha^2 &< \sigma_\beta^2 \leq \sigma_\gamma^2, \end{aligned}$$

where the relationship between σ_α^2 , σ_β^2 and σ_γ^2 follows logically if catch signals have been col-

lected from a , and sites a and b are correlated (if a is believed to have the highest abundance then mean abundance at b should be at least as large as mean abundance at the uncorrelated site c .) Based on the assumption that beliefs while mining a familiar site likely follow (9), we can use the model of section 2 to predict when a decision to gather information will occur.

Suppose that a skipper is mining a familiar site with beliefs characterized in (9). Repeat trips yield per trip catch signals, S_a . If these signals conform to expectations, $S_a = \alpha_0$, updated beliefs about mean abundance should not change, and the skipper should become more confident about the true abundance at a and at correlated site b . The skipper will continue to fish the familiar site; mining the familiar site should continue until catch signals from the site alter beliefs.

Now suppose catch at the familiar site falls below expectations, $S_a < \alpha_0$. From (2) and (3) the skipper lowers beliefs about α_0 and β_0 , while γ_0 is unchanged. Keeping in mind that all signals obtained reduce σ_α^2 and σ_β^2 , leaving σ_γ^2 unchanged, the incentive to gather information should be enhanced. First, as α_0 declines, the opportunity cost of gathering information at b and c declines. If exploration is warranted, where should it take place? >From (2)-(4) it is easy to see that a low signal at a lowers α_0 more than it lowers β_0 . This suggests updated beliefs following (I) of proposition 3. Combining this logic with the prediction of the proposition we conclude that if information acquisition occurs following a below-expectation catch signal from a familiar site it is likely to take place at an uncorrelated site.

Now suppose that a trip to familiar site a yields an above average signal. Updated beliefs based on good signals at site a cause α_0 and β_0 to increase, again leaving γ_0 unchanged (σ_α^2 and σ_β^2 decline and σ_γ^2 is unchanged). The updated beliefs that follow a good signal from a familiar site are most likely represented in (II) of proposition 3, with $\beta_0 \geq \tilde{\beta}_0$. This logic along with proposition 3 suggests that if information acquisition follows a good signal from a familiar site, it is likely to occur at a correlated site.

Lastly, if a fisherman chooses to gather information at correlated site b rather than independent site c , it must be the case that either $\sigma_\beta^2 > \sigma_\gamma^2$ or $\beta_0 > \gamma_0$. If a is familiar and b is believed positively spatially correlated, the most likely beliefs are $\sigma_\beta^2 < \sigma_\gamma^2$. Hence, if

information gathering occurs at a site that is positively correlated with a familiar site, the fisherman should believe that abundance at the correlated site is discretely larger than at uncorrelated sites, i.e., $\beta_0 > \gamma_0$.

3 Fishery background and data

The middle Atlantic surf clam fishery is located off the coast of New Jersey, Delaware, Maryland and Virginia. Surf clams are harvested year-round by towing dredges across underwater clam beds. Harvested clams are delivered to shore-based processing plants and consumed primarily in chowders and as fried clam strips. Surf clams are a long-lived, slow-growing bivalve species that reproduce by emitting spores that distribute on currents. Spores flourish at sites that provide the best habitat, e.g., bottom substrate and ample food supply. The habitat is heterogeneous and clams are concentrated in patches.

We obtained vessel- and trip-level logbook records from the National Marine Fisheries Service. The data include the port of departure, quantity of clams harvested, time at sea, time fishing, and the ten-minute square (measured using longitude and latitude) in which dredging occurred. Additional data from the Mid-Atlantic Fisheries Management Council, industry consultants, and other public sources include vessel characteristics, crew size and dockside clam prices. Data are available from 1983-2001.

Each ten-minute square (roughly 9 square miles) will represent a distinct fishing site. Using the latitude and longitude of a vessel's port and the midpoint of the ten-minute square as the site location, we are able to calculate steaming distances between the port and each site, as well as distances between sites.

Vessels typically range between 70-90 feet in length and are operated by a crew of 4-5. A typical trip involves steaming from port to a chosen site, dredging for roughly 6-10 hours and returning to port. Most trips are day trips. Our data indicate that preferred sites are located, on average, 50 miles from port. It is not unusual for a vessel to take 50-60 trips in a year.

	Mean	Std. Dev.	Minimum	Maximum
Vessel GRT	126.34	41.93	43	199
Vessel Length (ft.)	78.96	13.05	54	112
Crew Size	4.28	1.13	3	7
Trips	667.69	529.11	86	1914
Sites visited	39.93	17.24	8	69
Search trips per vessel (%)	13.50	8.03	2.81	33.72

Table 1: Vessel Descriptive Statistics and Site Visitation Patterns. GRT denotes gross registered tonnage. Searching trips are the first three trips to a previously unvisited site.

We treat each vessel as a distinct decision maker.⁶ To focus on sustained and representative search behavior, only vessels with a minimum 10 years of fishing activity are considered. The analysis that follows requires, to the extent possible, a complete temporal record of site visitation. Our data indicate that some vessels on some trips failed to record the latitude and longitude of a fished site. On occasion, inconsistent reporting of times at sea or dredging time was detected. Data errors of this sort make it difficult to distinguish new and previously visited sites, and to construct measures of performance at fished sites. Our empirical objective is to test site visitation patterns predicted by our model. We feel that this objective is best served with data containing minimal reporting error. Hence, we drop vessels with trip observations containing more than 1% anomalous observations.

Table 1 reports vessel characteristics and search activity for the 29 sample vessels. Vessels vary in size, as measured by gross registered tonnage (GRT), length, and crew size. The average number of trips taken by sample vessels is 667.69. Observations per vessel vary because some boats were active in the fishery longer than others. The fewest number of trip observations is 86 and the largest is 1,914. We construct a site- and date-specific visitation history for each vessel.⁷ Table 1 indicates that an average vessel visited 39.93 distinct sites

⁶It is possible that a sample vessel is operated by different captains throughout the data period. Industry sources indicate that captain turnover is infrequent.

⁷Fifteen vessels entered the surf clam fishery on or before 1983, the first year for which data are available. A history of sites visited prior to 1983 could not be created for these vessels. To avoid misrepresentation of search events, trips taken during the first two years (by vessels entering on or before 1983) are used to create a site visitation history only. Only 1985-2001 searching events are examined.

during the data period.

Clams are stationary; once a site with high abundance is found a skipper can make several return trips to *mine* the site's clams. Because clams do not move about, knowledge of the location of the most productive sites is valuable, i.e., information does not quickly decay.

The average length of a clam dredge is 110.68 inches. The area of a site exceeds 9 square miles, which is very large compared to the area that is scoured by the dredge. Our data indicate that 8 of the 29 vessels report taking 200 trips or more to a particular site. Vessels commonly take several trips to a new site before the site is abandoned (at least for the duration of our study period). We observe patterns where vessels visit a previously unfished site, then take one or more trips to a previously fished site, before returning. These search patterns, as well as discussions with industry, lead us to conclude that multiple trips to new sites can be considered part of a *search event*. Our data indicate that skippers frequently take up to 3 trips to a site before it is abandoned, whereas the frequency with which 4 or more trips are taken before a site is abandoned is markedly lower. This suggests that up to three signals are required before a skipper is sufficiently informed about clam abundance at a site. Based on the assumption the three trips to a new site can be classified as searching trips, Table 1 reports that the average vessel takes 13.50% of total trips searching new sites with considerable variation across vessels. The minimum and maximum search percentage is 2.81% and 33.72%, respectively.⁸

Clam fishermen have well over 200 distinct sites in which to dredge. It is well known among fishermen that stock abundance is positively spatially correlated. The patches in which clams are concentrated are small relative to the total size of the fishery. The ten-minute square sites are small relative to the size of these patches. Hence, as in the model of section 2, if fishermen decide to leave a familiar site and gather information, they may do so at a nearby, positively spatially correlated site or at a spatially distant, independent site.

Our data period spans a major shift in the management approach used to conserve the clam stock. Prior to October 1990, fleet harvest was regulated with limits on participating

⁸An extended appendix provides additional justification for the three trip cutoff. The results that follow were qualitatively indistinguishable under 3-5 trip cutoffs for search event signals.

vessels and limits on the hours per quarter that vessels could dredge for surf clams. After October 1990 input controls were replaced with an individual transferable quota (ITQ) program and fishing time restrictions were dropped.

4 Analysis and results

Our empirical goal is to investigate whether search and information acquisition by clam fishermen is consistent with the predictions of a model of rational search and learning under correlated information. For this purpose we impose minimal structure during the analysis.⁹ It should be emphasized that our approach is not immune to the identification problem that plagues empirical analysis of dynamic choice under uncertainty (see Manski, 2004 for a discussion). The results that follow are interpreted accordingly.

We first construct site-specific *signals* of clam stock abundance from the reported catch on each trip. Once at a chosen site, it is reasonable to assume that dredging proceeds until the expected marginal profit from additional dredging effort is zero. The catch thus equates expected marginal revenue product and marginal cost of effort. This suggests that differences in reported catch will be proportional to the sites perceived abundance. Also, sites located closer to port should be preferred to more distant sites. For each trip and each vessel we calculate a trip signal that is equal to the ratio of the observed harvest in bushels and the fuel required to steam to the site.

4.1 Do surf clam fishermen invest in information?

A first question to ask of the data is, can visits to new sites be characterized as investments in information? Finding that skippers invest in information about abundance at unfamiliar sites

⁹A structural discrete choice dynamic programming model would impose *a priori* assumptions for belief formation. A strategy for estimating a structural model in our setting is difficult to conceive. Nonstationary stock abundance, unobserved fishermen-specific choice sets, and correlated beliefs among others are complicating factors. Keane and Wolpin (1994) and Chintagunta et al. (2006) review these methods and their limitations.

would contradict myopic search behavior. We do not observe beliefs about true abundance and a direct test for myopic search is difficult to conceive. We instead examine catch signals to identify inconsistencies with myopic search. The following logic underlies the approach: If, on average, lower signals are obtained on search events than at familiar sites, it is reasonable to assume that fishermen expect to incur a cost in the form of a foregone expected catch while searching a new site, suggesting skippers do indeed make investments in information.

For each vessel we normalize the t 'th trip signal by a moving average of the signals obtained on 20 trips prior to trip t . We normalize to control for fishery-wide changes in clam abundance over time, information about abundance held by individual skippers, differences in vessel characteristics, and captain skill. The normalized catch signals are then pooled yielding 18,899 observations.

We next separate the normalized signals into three categories. The first includes signals obtained on search events; the data contain 1,856 search event signals. Search events yield signals that cause fishermen to update beliefs in one of two ways. If a favorable signal is received, the fisherman may return to mine the site's clams. If an unfavorable signal is obtained the site may be abandoned, at least in the short term. We define *mined* sites as sites that are eventually visited on more than 3 trips. The data contain 18,227 mined site signals. Lastly, *abandoned* sites are defined as those that are visited no more than 3 times. Our data contains 672 such search signals.¹⁰

Figure 2 depicts smoothed histograms for normalized signals received on all search events, search events at mined sites and search events at abandoned sites. An Epanechnikov kernel with bin width of 0.1 is used in constructing these histograms. We see that the distribution of signals obtained on search events is bimodal. This can be explained by differences in the beliefs that likely trigger search and will be investigated in later regressions. The sample mean value for signals obtained on search events is 0.942 with standard deviation 0.589. The sample mean of signals obtained at mined sites is 1.008 with standard deviation 0.380, and the sample mean of signals obtained at abandoned sites is 0.858 with standard deviation

¹⁰Catch signal categories are based on the visitation patterns of individual skippers. For example, a site may be classified as a mined site for one skipper and an abandoned site for another.

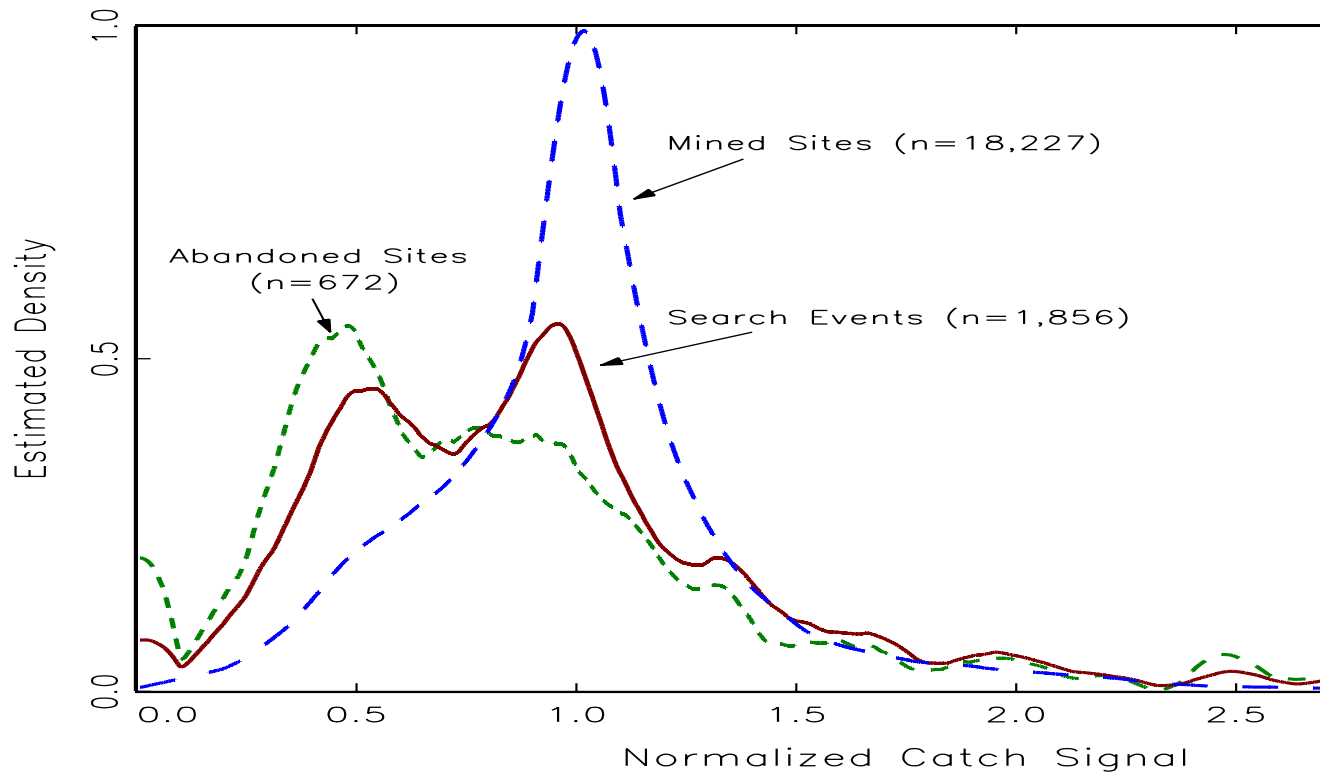


Figure 2: Signal Histogram at Searched, Mined and Abandoned Sites

0.662. A Mann-Wilcoxon-Whitney test of the null hypothesis that the signal distributions for search events and at mined sites are equivalent is rejected at the 99.5% level of confidence (the t-statistic is 13.279 with critical value 2.58). Figure 2 shows that the histogram of search event signals places more mass on low signals than the histogram for mined site signals. As expected, the signal distribution from abandoned sites has most mass on low signals. Notice that the relatively large support of the histogram of search event signals suggests that searching new sites is a likely risky activity.

Recalling that signals are normalized by a moving average, the results suggest that, on average, search events yield 94.2% of the reward obtained at sites that are fished prior to the search event. The implication is that when fishermen decide to search, they leave a site that on average yields a higher reward than the site chosen for search.

It is possible that clam fishermen’s actual beliefs about mean abundance at a new site are

biased upward; for example skipper may exhibit overconfidence bias regarding the quantity of clams they will discover when exploring a new site. If beliefs are unbiased however, they should be consistent with the observed signal distribution in our empirical data. In this case, results in Figure 2 would be consistent with the assertion that actual beliefs follow (9), and suggest that clam fishermen invest in information.

4.2 When and where do surf clam skippers gather information?

This section will investigate when and where information is gathered. The goal is to determine if the factors that trigger search in our model also trigger search in our data. We also consider if search location is consistent with the model predictions.

We begin by constructing a belief index that, under minimal assumptions, covaries with the skippers' unobserved beliefs about true clam stock abundance. A simple example demonstrates the construction of this index. Consider a skipper who takes a first trip to site j and obtains signal $S_{j,1}$. We make no assumptions for the expected catch on the first trip to a new site. The first trip signal $S_{j,1}$ is a proxy for the catch expectation for a second trip to site j . The second trip to site j yields a signal $S_{j,2}$ used to update beliefs in the event that a third trip to site j is taken, and so on. Implementing this procedure requires a belief updating rule to combine information gathered at each site. One possibility is naive updating where the catch expectation for the $t + 1$ trip to site j is equal to the signal obtained on the preceding trip. A second rule that we consider follows the Bayesian updating formulas from section 2. Under Bayesian updating, the expected catch on the t 'th visit is $\alpha_{j,t} = (S_{j,t-1}\sigma_{j,t-1}^2 + \alpha_{j,t-1}\sigma_s^2)/(\sigma_{j,t-1}^2 + \sigma_s^2)$, where $\sigma_{j,t-1}^2$ is the perceived variance at site j on trip $t - 1$.

Assuming that catch signals and actual beliefs covary, our constructed belief index tracks *changes* in actual beliefs. For example, a sequence of catch signals that exceed expectations will cause the belief index to exhibit a positive trend, whereas signals that fall short of expectations will cause the belief index to trend downward. To summarize changes in beliefs, we calculate the trend in the constructed belief index on the τ trips preceding trip t .¹¹ The

¹¹We found no qualitative changes in results for values of τ set equal to 5, 7 and 10. The results that

Variable	Model 1	Model 2
<i>Entrant</i>	0.32 (0.24)	0.37 (0.25)
<i>Price</i>	-0.53* (0.29)	-0.51* (0.29)
<i>Dredge time</i>	-2.49** (0.57)	-2.38*** (0.58)
ΔB_t	-0.59 (0.43)	-0.49 (0.40)
$(\Delta B_t)^2$	—	2.76** (1.37)
$E[\varphi]$	0.26 (0.54)	0.06 (0.54)
σ_φ^2	0.46*** (0.08)	0.45*** (0.08)
Yr	-1.42** (0.53)	-1.29** (0.53)
Yr^2	0.63** (0.23)	0.57** (0.23)
Log-likelihood	-3208.34	-3208.76

Table 2: Logistic Regression Results. Dependent variable is $d_{it} = 1$ if trip is first-time visit to a new site, zero otherwise. Single, double, and tripple asterisk indicates parameter is significant at the 90%, 95% and 99% level of confidence; ΔB_t - vessel-specific trend in constructed beliefs on preceding τ trips; *Dredge time* - indicator of time permitted to dredge for clams; *Price*- annual average dockside price; *Yr*- annual trend variable. There are 19,160 observations.

trend measure, which we denote as ΔB_t , is a proxy for the change in the fisherman’s beliefs about abundance at the sites that are fished prior to the t ’th trip. Additional details on the construction of the belief index and trend measure ΔB_t are presented in an extended appendix.

We construct a dependent variable d_{it} that is equal to 1 if trip t by vessel i is the first visit to a previously unfished site, and zero otherwise. The probability of that vessel i decides to search on trip t is $\Pr(d_{it} = 1) = F(\theta' z_{it} + \varphi_i)$, where F is the logistic cumulative distribution function, z_{it} is a vector of vessel i and trip t characteristics, θ is a parameter vector, and φ_i is a vessel-specific random effect representing unobserved preferences for search. We assume φ_i is distributed normally with mean $\bar{\varphi}$ and variance σ_φ^2 . The parameter θ is estimated using random effects logistic regression (McFadden and Train, 1997).

Table 2 reports parameter estimates, standard errors, and the log-likelihood values for

follow assume $\tau = 7$.

two model specifications. Each specification includes a linear and quadratic annual trend, Yr and Yr^2 , respectively, to control for unobserved changes in the true clam abundance during the data period. The parameters with annual trends, along with the mean and variance of the random vessel effects distribution are reported for completeness. These parameters are less important for understanding search behavior and are not discussed further.

Each specification in Table 2 includes a dummy variable, *Entrant*, set equal to 1 if a trip is taken within 1 year of the date the vessel entered the fishery, and 0 otherwise. New entrants have less information about abundance than experienced skippers, and likely search more. The results indicate that *Entrant* has a positive effect on the likelihood of search, although the result is not statistically significant.

The switch from controlled access management to rights-based management allows us to test whether the search behavior responded in expected ways to regulation. The variable *Dredge time* measures the hours per quarter that vessels were permitted to dredge.¹² It is set equal to 1 if no dredge time restrictions were in place and is otherwise equal to the permitted hours divided by the maximum permitted dredge hours in the data. Results indicate that the likelihood of search is inversely related to regulated dredging time. We interpret this result as an input substitution effect; by restricting dredging time only, the regulator lowered the opportunity cost of locating the highest abundance sites. Fishermen responded to the dredge time restriction by searching more. This argument is demonstrated formally in the appendix using a straight forward extension of our model.

The effects of risk aversion on the *demand* for information are ambiguous (Freixas and Kihlstrom, 1984). A visit to an unfamiliar site exposes the fisherman to risk; the cost of gathering this information is higher for a risk averse fisherman. However, signals from unfamiliar sites can reduce uncertainty on subsequent trips. The net effect on the demand for information is ambiguous.¹³

We will use the following approach to analyze the risk preferences of clam fishermen.

¹²Dredging time restrictions began in 1984 and ended when individual transferable quotas were adopted in October, 1990. In 1988, vessels were restricted to fishing 6 hours every other week.

¹³By invoking a continuity of preferences argument, the results that we obtain for a risk neutral fisherman carry through under risk averse preferences.

The risk on a fishing trip derives primarily from revenue risk, i.e., dockside price times the uncertain per trip catch. The variable *Price* is the real annual average dockside price of surf clams. A price increase raises the perceived revenue variance at the rate $Price^2$, times the perceived catch variance. As the price rises, revenue variance at an unfamiliar site increases at a faster rate than revenue variance at familiar sites. Importantly however, the information about true stock abundance that is contained in the catch signal is unaffected by the dockside price. For a risk averse fisherman, the cost of information, in terms of the disutility from bearing risk, rises with the *Price* variable. All else equal, the amount of search will decrease with *Price* if clam fishermen are risk averse, whereas search will increase (remain unaffected) under risk loving (risk neutral) preferences. The results in Table 2 find that *Price* has a negative and significant effect on the likelihood of search indicating risk averse preferences. We will return to the issue of risk aversion below.

Model 1 of Table 2 finds that the probability of a search event is inversely related to the trend measure ΔB_t . The estimated standard error indicates that the effect is not significantly distinguishable from zero. To test whether search events are also more likely when ΔB_t is positive, Model 2 of Table 2 includes the regressor $(\Delta B_t)^2$. The results find that the effect of the linear term on search remains negative while the quadratic term has a positive and significant effect on the likelihood of search. The likelihood of a search event is found to increase with largely negative and largely positive values of ΔB_t . The fitted parameters indicate that the likelihood of a search event is smallest when $\Delta B_t = 0.054$, with a 95% confidence interval that includes zero (using the delta method, we calculate this confidence interval at $[-0.103, 0.212]$). Figure 3 shows the fitted probability of search as a function of the trend measure ΔB_t .

Two explanations for the increased search when ΔB_t is largely positive come to mind. One is that fishermen's utility satisfies decreasing absolute risk aversion, so that a string of productive trips, which will be the case if $\Delta B_t > 0$, increases willingness to bear risk. This explanation is somewhat contrary to the finding that less search occurs when prices are high. Higher prices mean more income and should increase search if preferences satisfy decreasing absolute risk aversion.




Figure 3: Search probability conditional on ΔB_t

An alternative explanation, one that is consistent with the model of section 2, is that favorable signals obtained at a familiar site, i.e., the case of $\Delta B_t > 0$, cause fishermen to update beliefs about abundance at the site yielding the signals and at correlated or spatially nearby sites. Based on the updated beliefs, skippers make a trip to a correlated site to gather information. We are able to investigate this hypothesis by examining the relationship between the variable ΔB_t and the location of search.

We calculate the distance between sites that are selected for search and sites visited on the τ trips preceding a search event. For each of the 770 first-time visits to a new site we calculate the distance between the new site and the vessel's most productive sites on the previous τ trips. This distance is then regressed on the change-in-belief variable ΔB_t , the annual trend variables Yr and Yr^2 , and vessel-specific effects φ_i . Ordinary least squares finds that the average distance between a new site and recently fished sites is 37.96 miles. The marginal

effect of ΔB_t on this distance is -5.68 miles, with a 95% confidence interval [-7.85, -3.51]. This result suggests that search events conditional on increasing beliefs about abundance at familiar sites occur near the sites yielding the favorable signals, whereas search events conditional on declining beliefs about abundance at familiar sites occur at spatially distant sites. This search pattern is consistent with our model of rational search with correlated information.

As a final check on the consistency of this search pattern, we re-examine the pooled sample of normalized signals presented in Figure 2. We have argued that with $\Delta B_t > 0$, information gathering is likely to occur at a correlated site rather than an uncorrelated site since a good catch at a familiar site raises expectations at correlated sites. This suggests that search event signals conditional on $\Delta B_t > 0$ should, on average, exceed search event signals conditional on $\Delta B_t < 0$. Of the 1,856 search event signals, 1,144 occurred with $\Delta B_t < 0$, and 712 occurred with $\Delta B_t > 0$. Conditional on $\Delta B_t < 0$, the sample average signal is 0.926, with standard deviation 0.520. Conditional on $\Delta B_t > 0$, the sample average signal is 0.967, with standard deviation 0.684. Hence, the bimodal distribution of search event signals in Figure 2 adds further empirical support to rational search with correlated information as an explanation for increased search when $\Delta B_t > 0$.

The above analyses utilize a belief index whose construction mirrors a Bayesian updating rule. A comprehensive test of alternate updating rules is beyond the scope of this study. We can compare the results so far with a belief index that is constructed following naive belief formation. For this test we recalculate our change-in-beliefs index based on the *signals* obtained on the τ trips that precede each trip. This generates a change-in-beliefs measure that is based on a naive updating rule. We repeated each regression in Table 2 using the naive change-in-belief variable. Based on a comparison of likelihood function values, we reject this model in favor of the index ΔB_t . This finding does not allow us to conclude clam fishermen are Bayesian. It does suggest that a naive updating rule provides an inferior fit to the search patterns in our data.

5 Conclusion

This paper studies search and information acquisition in the presence of correlated learning. A simple model is used to investigate the effects of correlation on the decision to acquire information, and the source of the information. Using an empirical approach that relies of minimal assumptions for belief formation and learning, we find that the site choices, and particularly the decisions by fishermen to acquire information about the true location of surf clam stocks, are consistent with the predictions of our model. We find evidence that suggests surf clam skippers regularly incur a cost in the form of a foregone catch at a familiar site, to gather information about abundance at new sites. Clam fishermen are more likely to search when the trend in catch performance at recently fished (familiar) sites is largely negative and largely positive. When the catch at familiar sites declines, so does the cost of gathering information. When the catch at recently fished sites is on the rise, skippers should update beliefs about true abundance at the recently fished site and at sites believed to have positively correlated abundance. Results indicate that search is more likely when catch at recently fished sites increases. Moreover, we find that when search follows a sequence of productive trips it tends to occur at nearby sites whereas search that follows a sequence of low productivity trips tends to occur at distant sites, i.e., sites that are uncorrelated with the low-productivity site.

Overall our results suggest that professional clam fishermen balance exploitation and exploration in ways that are consistent with the predictions of a model of rational search and learning. Due to identification limitations in models of dynamic choice under uncertainty however a formal test of rational search and learning is not possible. We cannot conclude that clam fishermen rationally search and learn. Alternative learning rules, for example, rules that give arbitrary positive weight to past and new catch signals when forming beliefs about catch success, would also be consistent with our data, as would a directed cognition model of learning (Gabaix and Laibson, 2000). However, we do find that myopic search and naive belief updating rules are less able to explain the search patterns observed in our data.

The finding that fishing site choice patterns appear consistent with a model of correlated information and learning has implications for general problems of search under uncertainty.

Analysis of search and learning typically assumes that payoffs at available search locations are independent. The independence assumption can dramatically simplify analysis, but may not capture important aspects of reality (see also Gabaix and Laibson, 2000). Correlated information is a natural characteristic of many search problems. For example, with directed job search, an offer from a firm may provide information about the wage distribution at the firm's competitors. Searching for valuable genetic material in a plant likely provides information about the prospects of finding valuable materials in genetically related plants. Similar examples could be constructed for clinical trials, research and development, oil and gas exploration and a host of other situations. Further research into learning with correlated information could refine our understanding of search in a variety of settings and lead to more predictive models.

Lastly, our findings have practical implications for fisheries management and the spatial movements of fishermen. Our finding that clam fishermen searched more in response to dredge time restrictions may have undermined the intent of the regulation. Managers report that landings per unit of effort tripled during the period that dredge time restrictions were used (Mid-Atlantic Fisheries Management Council, 1999). While increases in clam stock abundance could explain this sharp increase, it may have resulted from increased search and an increase in the stock of knowledge about the true location of clams. Studies of the spatial movements of fishermen typically assume, for simplicity, that site choices are made by myopic decision makers who use naive updating rules to form catch expectations.¹⁴ Models that acknowledge active exploration, consistent with forward-looking agents, and less restrictive learning rules could improve the understanding of the spatial movements of fishermen and in turn, improve the management of fisheries (Sanchirico and Wilen, 2005).

¹⁴Smith (2000) reviews the spatial search literature in fisheries. Mangel and Clark, 1983, study rational search by a fishing fleet, however, they do not present a formal analysis site choice.

6 References

- Akerberg, D. A.** “Empirically distinguishing informative and prestige effects of advertising,” *Rand Journal of Economics* **32** (2001): 316-333.
- Adam, K.** “Learning while searching for the best alternative,” *Journal of Economic Theory* **101** (2001): 252-280.
- Banks, J., M. Olsen and D. Porter.** “An experimental analysis of the bandit problem,” *Economic Theory* **10** (1997): 55-77.
- Ching A. T.** “Consumer learning and heterogeneity: dynamics of demand for prescription drugs after patent expiration,” mimeo, University of Toronto, 2005.
- Chintagunta, P., T. Erdem, P. E. Rossi and M. Wedel** “Structural Modeling in Marketing: Review and Assessment” *Marketing Science* **25** (2006): 604-616.
- Crawford G. S., and M. Shum.** “Uncertainty and learning in pharmaceutical demand,” *Econometrica* **73** (2005): 1137-1173.
- De Groot, M. H.** *Optimal Statistical Decisions*, McGraw-Hill, New York, 1970.
- Erdem, T. and M. P. Keane.** “Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets,” *Marketing Science* **15** (1996): 1-20.
- Erdem, T. M. P. Keane and J. Strebel.** “Learning about computers: An analysis of information search and technology choice,” *Quantitative Marketing and Economics* **3** (2005): 207-247.
- Freixas, X., and R. E. Kihlstrom.** “Risk aversion and information demand,” in *Bayesian Models in Economic Theory*, M. Boyer and R. E. Kihlstrom eds., Elsevier, New York, NY, 1984.
- Garicano, L.** “Hierarchies and the organization of knowledge in production,” *Journal of Political Economy* **108** (2000): 874-904.
- Gabaix X. and D. Laibson.** “Bounded rationality and directed cognition,” mimeo, Massachusetts Institute of Technology and Harvard University, 2000.

- Gittins, J. C.** “Bandit processes and dynamic allocation indices,” *Journal of the Royal Statistical Society B* **41** (1979): 148-177.
- Keane, M. P., and K. I. Wolpin.** “The solution and estimation of discrete dynamic programming models by simulation and interpolation: Monte Carlo evidence,” *The Review of Economics and Statistics* **76** (1994): 648-672.
- Mangel, M. and C. Clark.** “Uncertainty, search and information in fisheries,” *Journal of the International Council for the Exploration of the Seas* **41** (1983): 93-103.
- Manski, C. F.** “Measuring expectations,” *Econometrica* **72** (2004): 1329-1376.
- McFadden, D., and K. Train.** “Mixed multinomial logit models for discrete responses,” *Journal of Applied Econometrics* **15** (2000): 447-470.
- Meyer, R. J., and Y. Shi.** “Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem,” *Management Science* **41** (1995): 817-834.
- Mid-Atlantic Fisheries Management Council.** “Overview of the Surfclam and Ocean Quahog Fisheries and Quota Recommendations for 2000,” Dover, DE, August, 1999.
- Sanchirico, J., and J. Wilen.** “Optimal spatial management of renewable resources: matching policy scope to ecosystem scale,” *Journal of Environmental Economics and Management* **50** (2005): 23-46.
- Smith, M.** “Spatial search and fishing location choice: Methodological challenges of empirical modeling,” *American Journal of Agricultural Economics* **82** (2000): 1198-1206.
- Sulganik E., and I. Zilcha.** “The value of information: The case of signal-dependent opportunity sets,” *Journal of Economic Dynamics and Control* **21** (1997): 1615-1625.
- Zellner, A.** *An Introduction to Bayesian Inference in Econometrics*, John Wiley and Sons, Inc., New York, 1971.

7 Appendix

7.1 Calculations and proposition proofs

Fishing at site b in period one yields S_b , leading to updated beliefs:

$$\begin{aligned} (\beta|S_b) &\sim N\left(\delta\frac{S_b\sigma_\beta^2 + \beta_0\sigma_s^2}{\sigma_\beta^2 + \sigma_s^2}, \delta^2\frac{\sigma_\beta^2\sigma_s^2}{\sigma_\beta^2 + \sigma_s^2}\right) \\ (\alpha|S_b) &\sim N\left(\alpha_0 + \rho\frac{\sigma_\alpha\sigma_\beta}{(\sigma_\beta^2 + \sigma_s^2)}(S_b - \beta_0), \sigma_\alpha^2\left(1 - \rho^2\frac{\sigma_\beta^2}{(\sigma_\beta^2 + \sigma_s^2)}\right)\right). \end{aligned}$$

Solution for \underline{S}_a and \underline{S}_b :

$$\begin{aligned} \underline{S}_a &= \beta_0\frac{\sigma_\alpha^2 + \sigma_s^2}{\delta\sigma_\alpha^2 - \sigma_{\alpha\beta}} - \alpha_0\frac{\sigma_{\alpha\beta} + \delta\sigma_s^2}{\delta\sigma_\alpha^2 - \sigma_{\alpha\beta}} \\ \underline{S}_b &= \alpha_0\frac{\sigma_\beta^2 + \sigma_s^2}{\delta\sigma_\beta^2 - \sigma_{\alpha\beta}} - \beta_0\frac{\sigma_{\alpha\beta} + \delta\sigma_s^2}{\delta\sigma_\beta^2 - \sigma_{\alpha\beta}}. \end{aligned}$$

Assuming $\delta = 1$, the standardized threshold signals $\underline{\lambda}_a$ and $\underline{\lambda}_b$ are,

$$\begin{aligned} \underline{\lambda}_a &= \frac{\underline{S}_a - \alpha_0}{\sqrt{\sigma_\alpha^2 + \sigma_s^2}} = -(\alpha_0 - \beta_0)\frac{\sqrt{\sigma_\alpha^2 + \sigma_s^2}}{\sigma_\alpha^2 - \sigma_{\alpha\beta}} < 0 \\ \underline{\lambda}_b &= \frac{\underline{S}_b - \beta_0}{\sqrt{\sigma_\beta^2 + \sigma_s^2}} = (\alpha_0 - \beta_0)\frac{\sqrt{\sigma_\beta^2 + \sigma_s^2}}{\sigma_\beta^2 - \sigma_{\alpha\beta}} > 0. \end{aligned}$$

Proof of proposition 1: The proof goes as follows. After establishing analytical expressions for $V(b)$, $V(a)$ and $\Delta V = V(b) - V(a)$, we prove uniqueness by showing monotonicity of ΔV . Existence is proven by showing, together with the continuity of the function $\Delta V(\cdot)$, that ΔV changes sign when $\sigma_\beta^2 \in (0, +\infty)$.

Recall the expressions for $V(b)$ and $V(a)$:

$$V(b) = \beta_0 + [1 - \Phi(\underline{\lambda}_b)]E[\beta|S_b \geq \underline{S}_b] + \Phi(\underline{\lambda}_b)E[\alpha|S_b < \underline{S}_b]$$

and

$$V(a) = \alpha_0 + [1 - \Phi(\lambda_a)]E[\alpha|S_a \geq \underline{S}_a] + \Phi(\lambda_a)E[\beta|S_a < \underline{S}_a],$$

where

$$E[\beta|S_b \geq \underline{S}_b] = \frac{S_b^+ \sigma_\beta^2 + \beta_0 \sigma_s^2}{\sigma_\beta^2 + \sigma_s^2}$$

$$E[\alpha|S_a \geq \underline{S}_a] = \frac{S_a^+ \sigma_\alpha^2 + \alpha_0 \sigma_s^2}{\sigma_\alpha^2 + \sigma_s^2},$$

and

$$E[\alpha|S_b < \underline{S}_b] = \alpha_0 + \frac{\sigma_{\alpha\beta}}{(\sigma_\beta^2 + \sigma_s^2)}(S_b^- - \alpha_0)$$

$$E[\beta|S_a < \underline{S}_a] = \beta_0 + \frac{\sigma_{\alpha\beta}}{(\sigma_\alpha^2 + \sigma_s^2)}(S_a^- - \alpha_0),$$

where S_b^+ , S_b^- , S_a^- , and S_a^+ denote truncated mean signals as defined in section 2.V (b) and $V(a)$ may be written as

$$V(b) = 2\beta_0 + (\alpha_0 - \beta_0) \Phi(\lambda_b) + \phi(\lambda_b) \frac{\sigma_\beta^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\beta^2 + \sigma_s^2}}$$

$$V(a) = 2\alpha_0 - (\alpha_0 - \beta_0) \Phi(\lambda_a) + \phi(\lambda_a) \frac{\sigma_\alpha^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\alpha^2 + \sigma_s^2}}.$$

$\Delta V = V(b) - V(a)$ will denote the relative value of investing in information.

We have

$$\Delta V = -(\alpha_0 - \beta_0) [2 - \Phi(\lambda_b) - \Phi(\lambda_a)] + \phi(\lambda_b) \frac{\sigma_\beta^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\beta^2 + \sigma_s^2}} - \phi(\lambda_a) \frac{\sigma_\alpha^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\alpha^2 + \sigma_s^2}}.$$

Using the relationship

$$\frac{\partial \phi(\lambda_b)}{\partial \sigma_\beta^2} = -\lambda_b \frac{\partial \lambda_b}{\partial \sigma_\beta^2} \phi(\lambda_b),$$

the derivative of ΔV with respect to σ_β^2 is computed as

$$\frac{\partial \Delta V}{\partial \sigma_\beta^2} = (\alpha_0 - \beta_0) \phi(\underline{\lambda}_b) \frac{\partial \underline{\lambda}_b}{\partial \sigma_\beta^2} - \underline{\lambda}_b \frac{\partial \underline{\lambda}_b}{\partial \sigma_\beta^2} \phi(\underline{\lambda}_b) \frac{\sigma_\beta^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\beta^2 + \sigma_s^2}} + \frac{(\sigma_\beta^2 + \sigma_{\alpha\beta} + 2\sigma_s^2)}{2(\sigma_\beta^2 + \sigma_s^2) \sqrt{\sigma_\beta^2 + \sigma_s^2}} \phi(\underline{\lambda}_b).$$

Finally, replacing $\underline{\lambda}_b$ and $\partial \underline{\lambda}_b / \partial \sigma_\beta^2$ by their respective values, we obtain after simplification

$$\frac{\partial \Delta V}{\partial \sigma_\beta^2} = \phi(\underline{\lambda}_b) \left(\frac{\sigma_\beta^2 + \sigma_{\alpha\beta} + 2\sigma_s^2}{2(\sigma_\beta^2 + \sigma_s^2) \sqrt{\sigma_\beta^2 + \sigma_s^2}} \right) > 0.$$

This proves uniqueness. To prove existence, first recall that $\Delta V(\cdot)$ is continuous on $(0, +\infty)$.

Then notice that when $\sigma_\beta^2 = \sigma_\alpha^2$, $\underline{\lambda}_a = -\underline{\lambda}_b$, thus ΔV becomes

$$\Delta V = V(b) - V(a) = -2(\alpha_0 - \beta_0) + (\alpha_0 - \beta_0) (\Phi(\underline{\lambda}_b) - \Phi(-\underline{\lambda}_b)) < 0.$$

When $\sigma_\beta^2 \rightarrow \infty$, $V(a)$ has a finite value. Moreover, we also check that $\underline{\lambda}_b \rightarrow 0$ when $\sigma_\beta^2 \rightarrow \infty$.

Therefore,

$$\lim_{\sigma_\beta^2 \rightarrow \infty} V(b) = 2\beta_0 + \Phi(0)(\alpha_0 - \beta_0) + \phi(0) \lim_{\sigma_\beta^2 \rightarrow \infty} \left(\frac{\sigma_\beta^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\beta^2 + \sigma_s^2}} \right) > 0.$$

This concludes the proof of the existence of a unique $\bar{\sigma}_\beta^2$ such that $\Delta V = 0$.

Proof of Proposition 2: To prove this result, we compute the derivative of ΔV with respect to $\sigma_{\alpha\beta}$. Keeping in mind that

$$\frac{\partial \phi(\underline{\lambda}_b)}{\partial \sigma_{\alpha\beta}} = -\underline{\lambda}_b \frac{\partial \underline{\lambda}_b}{\partial \sigma_{\alpha\beta}} \phi(\underline{\lambda}_b) \quad \text{and} \quad \frac{\partial \phi(\underline{\lambda}_a)}{\partial \sigma_{\alpha\beta}} = -\underline{\lambda}_a \frac{\partial \underline{\lambda}_a}{\partial \sigma_{\alpha\beta}} \phi(\underline{\lambda}_a),$$

we obtain

$$\begin{aligned} \frac{\partial \Delta V}{\partial \sigma_{\alpha\beta}} &= (\alpha_0 - \beta_0) \left(\phi(\underline{\lambda}_b) \frac{\partial \underline{\lambda}_b}{\partial \sigma_{\alpha\beta}} + \phi(\underline{\lambda}_a) \frac{\partial \underline{\lambda}_a}{\partial \sigma_{\alpha\beta}} \right) + \\ &- \phi(\underline{\lambda}_b) \left(\underline{\lambda}_b \frac{\partial \underline{\lambda}_b}{\partial \sigma_{\alpha\beta}} \frac{\sigma_\beta^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\beta^2 + \sigma_s^2}} + \frac{\sigma_{\alpha\beta}}{\sqrt{\sigma_\beta^2 + \sigma_s^2}} \right) + \phi(\underline{\lambda}_a) \left(\underline{\lambda}_a \frac{\partial \underline{\lambda}_a}{\partial \sigma_{\alpha\beta}} \frac{\sigma_\alpha^2 - \sigma_{\alpha\beta}}{\sqrt{\sigma_\alpha^2 + \sigma_s^2}} + \frac{\sigma_{\alpha\beta}}{\sqrt{\sigma_\alpha^2 + \sigma_s^2}} \right). \end{aligned}$$

Using the analytical expressions of $\underline{\lambda}_b$ and $\underline{\lambda}_a$, basic algebra shows that the expression above

simplifies to

$$\frac{\partial \Delta V}{\partial \sigma_{\alpha\beta}} = \sigma_{\alpha\beta} \left(\frac{\phi(\lambda_a)}{\sqrt{\sigma_\alpha^2 + \sigma_s^2}} - \frac{\phi(\lambda_b)}{\sqrt{\sigma_\beta^2 + \sigma_s^2}} \right).$$

Thus, we have $\frac{\partial \Delta V}{\partial \sigma_{\alpha\beta}} \leq 0$ if

$$\frac{\phi(\lambda_a)}{\phi(\lambda_b)} \leq 1 \leq \frac{\sqrt{\sigma_\alpha^2 + \sigma_s^2}}{\sqrt{\sigma_\beta^2 + \sigma_s^2}}$$

holds true for any $\sigma_\beta^2 \geq \sigma_\alpha^2$. Note first that $\lambda_a = -\lambda_b$ when $\sigma_\beta^2 = \sigma_\alpha^2$ and, using the symmetry property of $\phi(\cdot)$, the (double) inequality reduces to an equality. When $\sigma_\beta^2 > \sigma_\alpha^2$, the right-hand side inequality holds strictly. To show that the left-hand side holds strictly as well, first recall that $\partial \lambda_b / \partial \sigma_\beta^2 < 0$ and then compute

$$\frac{\partial \left(\frac{\phi(\lambda_a)}{\phi(\lambda_b)} \right)}{\partial \sigma_\beta^2} = \lambda_b \frac{\partial \lambda_b}{\partial \sigma_\beta^2} \frac{\phi(\lambda_a)}{\phi(\lambda_b)} < 0 \text{ for any } \sigma_\beta^2 > \sigma_\alpha^2.$$

This concludes the proof.

7.2 Dredge time restrictions and the incentive to search

We want to show that a regulation that restricts dredging time sharpens the incentive to search. In the baseline model of section 2 there is no dredge restriction and payoffs are collected during two fishing periods. Now assume that the restriction allows fisherman to dredge a maximum of 1 time unit in each period and a total $(1 + k)$ units in two periods with $k < 1$. We have the following result:

Lemma If the fisherman is indifferent between a search event and no search event when no dredging time restriction exist, then he strictly prefers to search under a dredge time restriction, i.e., when $k < 1$.

A *dr* superscript will distinguish the case with a dredge time restriction. The expected payoff when the fisherman chooses to visit site *a* first (i.e., chooses not to explore the uncertain

site b) is summarized as

$$V(a) = \alpha_0 + \underbrace{\max [\text{return to } a, \text{ switch to } c]}_A.$$

Note that the second right-hand term, A , is the period 2 expected payoff under an optimal site choice strategy. The expected catch when the fisherman chooses to search, i.e., visit b first is

$$V(b) = \beta_0 + \underbrace{\max [\text{return to } a, \text{ switch to } b, \text{ switch to } c]}_B.$$

If the fisherman is indifferent between search and no search, we have $V(a) = V(b)$, which implies

$$(10) \quad \alpha_0 - \beta_0 = B - A$$

and

$$\alpha_0 < A \text{ and } \beta_0 < B.$$

Now consider the dredge restriction. To fully benefit from information gathered, the fisherman optimally allocates 1 unit of dredge time to the second period and k units to the first period (the exploratory period). The expected returns with a dredge restriction are

$$V^{dr}(a) = k\alpha_0 + A \text{ and } V^{dr}(b) = k\beta_0 + B.$$

The fisherman strictly prefers to search if $V^{dr}(b) > V^{dr}(a)$ or if $B - A > k(\alpha_0 - \beta_0)$. This will be the case given (10) and $k < 1$.


8 Extended Appendix

This extended appendix accompanies “Search and active learning with correlated information: Empirical evidence from Mid-Atlantic clam fishermen.” The first section provides additional justification for the assumption that the first three signals obtained at a new site are search event signals. The following section describes the calculation of the belief index and trend measure used in the paper.

8.1 Search event trips

Our empirical procedures require an assumption for the number of signals that can be assumed search signals. We examine the data, and in particular the abandonment of fishing sites by the clam skippers. For each fisherman in our sample we determine the unique sites that were visited. We then count the number of visits to each of these sites. The following histogram shows the frequency with which a fishermen visited a given site exactly k times, $k = 1, 2, 3, \dots$. Note that the sites that are visited more than 20 times are not included.

The idea is that if a site is explored once and then abandoned, one might assume that the fisherman needed only one signal to learn about the site’s true stock abundance. In this case, one might assume that only one trip to a new site is required to learn about the site’s true stock abundance.



Percentage of sites that are visited $k = 1, 2, 3, \dots, 20$ times.

The data reveal a substantial drop in the frequency of visits following 3 trips. This suggests that clam fishermen frequently take up to 3 trips to a site before it is abandoned. The frequency with which they take 4 trips to the site before it is abandoned is markedly less. A similar, although less pronounced, drop in the frequency occurs between 5 and 6 trips to a site. Note that 27.7% of the sites received more than 20 visits (this frequency is not shown in the above figure). These are sites from which high catch signals were presumably obtained and were eventually mined by the fishermen in our data.

This analysis of site abandonment patterns supports a cutoff somewhere in the 3-5 trip range. Distinguishing further between the 3, 4, or 5 trip cutoff is a difficult matter. We analyzed our data under the 3, 4 and 5 trip cutoff and find that the results are remarkably robust.

8.2 Belief index and trend measure

An example illustrates the procedure used to construct our index. The table below depicts a short sequence of trips taken by a mock fisherman. The table reports the number of trips (8 in all) selected sites 1, 2 or 3, signals received, updated beliefs and the trip type. Note that

in this example only first time visits to a site are considered search events and all other trips are considered mining trips. We have assumed uncorrelated beliefs across the three sites.

Trip number	Selected Site	Signal Obtained	Site 1 beliefs	Site 2 beliefs	Site 3 beliefs	Trip Type
1	1	$S_{1,1}$	—	—	—	search event
2	1	$S_{1,2}$	$\alpha_{1,2}$	—	—	mining trip
3	1	$S_{1,3}$	$\alpha_{1,3}$	—	—	mining trip
4	2	$S_{2,1}$	$\alpha_{1,3}$	—	—	search event
5	2	$S_{2,2}$	$\alpha_{1,3}$	$\alpha_{2,2}$	—	mining trip
7	3	$S_{3,1}$	$\alpha_{1,3}$	$\alpha_{2,2}$	—	search event
8	3	$S_{3,2}$	$\alpha_{1,3}$	$\alpha_{2,2}$	$\alpha_{3,2}$	mining trip

The mock fisherman visits site 1 on the first trip and obtains signal $S_{1,1}$. This signal is used to update beliefs about the expected abundance at site 1, denoted $\alpha_{1,2}$. The first trip is a search event since no sites have been previously fished. The second trip, also to site 1, yields signal $S_{1,2}$, which is used in forming belief $\alpha_{1,3}$, the expected abundance on the third trip to site 1. Similarly, the third trip signal, $S_{1,3}$, is used to update beliefs in the event that a fourth trip is taken to site 1. In the example, the fourth trip to site 1 does not occur. Assuming no information decay, the mock fisherman's beliefs about expected abundance at site 1 are fixed at $\alpha_{1,3}$ for the remainder of the data period. A similar process is used to construct beliefs about expected abundance at sites 2 and 3.


Under the assumptions in the paper, the expected catch on the second trip to site j is $\alpha_{j,2} = S_{j,1}$. The expected catch for the third trip to site j is then $\alpha_{j,3} = (S_{j,2}\sigma_{j,2}^2 + \alpha_{j,2}\sigma_s^2)/(\sigma_{j,2}^2 + \sigma_s^2)$, where $\sigma_{j,2}^2$ is the updated variance at site j following the second trip and σ_s^2 denotes signal variance. The expected catch on the t 'th visit to the site is $\alpha_{j,t} = (S_{j,t-1}\sigma_{j,t-1}^2 + \alpha_{j,t-1}\sigma_s^2)/(\sigma_{j,t-1}^2 + \sigma_s^2)$.

To implement this procedure with data we must specify a value for the signal variance σ_s^2 , and a value for the perceived variance of abundance on the first visit to a new site. An estimate of the signal variance is obtained as the signal variance at mined sites. The idea is that signal variation while mining at familiar sites represents usual trip to trip variation in

catch success. Following similar logic, the perceived variance of abundance at a new site is assumed equal to the signal variance on search events. Search event signal variance is roughly 40% higher than mined site signal variance. Note that these variance parameters determine the weight that is placed on past and current signals in the belief updating process. We examine the sensitivity of the results to modest changes in these parameter values and find only minor differences.

Under the assumption that catch signals and actual beliefs covary, the constructed beliefs track *changes* in actual beliefs. For example, a sequence of signals that exceed expectations will cause the belief index to exhibit positive trend, whereas recent signals that conform to expectations will cause no trend in the constructed beliefs.

To summarize changes in beliefs, which are key for predicting when a search event will occur, we measure the trend in the constructed belief index for the τ trips preceding trip t . Denote the vessel-specific trip t trend measure as ΔB_t .



Trend in constructed beliefs

The above figure illustrates how the ΔB_t measure is obtained. The diagram shows (mock) constructed beliefs on τ trips preceding trip t . The beliefs trend downward and thus for trip t , ΔB_t is shown to be negative. The trend measure is calculated by regressing the constructed beliefs on trip $t - \tau - 1$ through $t - 1$ against a constant and linear trend variable. ΔB_t is then set equal to the slope parameter from this regression.