

1. Introduction

We calibrate and validate the POSEIDON (Bailey et al., 2018) fisheries agent-based model using data from the US West Coast groundfish fishery and compare the performance of simple, adaptive algorithms with imperfect information with other, more commonly used decision-making algorithms that include perfect information and/or rationality. We show that the adaptive algorithms explain observed data better. Moreover, while it is possible to derive statistical agents from logbook data, a simple, adaptive, uncalibrated decision-making algorithm performs out-of-sample just as accurately.

We address two gaps in the literature. First, we compare the usual assumption in bioeconomic models of allocating effort automatically where profits would be maximized ("assuming away the problem of finding fish", Wilson, 1990) to more "bounded" rationality, either in terms of information available or ability to process it. Second, we implement multiple decision-making processes within the same bioeconomic model; this way decision-making algorithms can be compared not just by their ability to predict future actions but also on the system-wide effects they have over the biology and economic performance.

Two recent bioeconomic models focused on the US West Coast groundfish fishery. Toft et al. (2011) modelled groundfish trawlers as they entered the individual transferable quota (ITQ) program while Kaplan et al (2014) modelled the effects of 20 fleets,

representing gear types, on the whole California Current ecosystem. Both studies assume agents know perfectly the profits they will make in each area before making a trip.

More generally, Van Putten et al (2012) classifies behavioural models of fishers into three groups: dynamic optimization, discrete-choice models or agent-based models. For all three it is rare to model exploration and learning directly; more common is to either assume “perfect knowledge”, agents knowing already the profits they will make before travelling, or rational expectations, agents having the correct expectations of what profits or catches will be.

Dynamic optimization fishers compute the optimal long term plan by allocating effort in time and space by value iteration (Clark & Mangel, 2000). Because of its computational complexity (see Littman, Dean & Kaelbling, 1995) perfect knowledge is an important expedient to keep the problem dimension small and computable. Dynamic programming fishers in Dowling et al. (2012) not only know abundance throughout the ocean but also stock dynamics and migratory patterns. Similarly, the Alaskan multispecies groundfish trawlers in Ono et al (2017) know the yields of all *metiers*, mixing them optimally through linear programming. Boettiger, Mangel & Munch (2015) manage to add uncertainty to a dynamic programming problem but to do so they remove geography and have only one representative fisher (which makes it hard to study allocative results of policies).

Expectations of catches or profits are a key component in discrete-choice models.

These expectations can be the actual catches or profits the fisher will make or a noisy, lagged observation of them as in Mistiaen and Strand (2000). Fishers in Haynie and Layton (2010) know correctly the average catches they will make in each area as do the recreational fishers in Baerenklau and Provencher (2005). In the dynamic model of Hicks and Schnier (2006), fishers not only perfectly predict catches in every area but can mentally simulate their evolution through time and no new information can be obtained by either searching or fishing.

Not all statistical discrete-choice models abstract away from imperfect knowledge and exploration, however. Abbot & Wilen (2011) abandon the usual assumptions of shared knowledge among the whole fleet by modelling explicitly information sharing among boats. Agents in Hutniczak & Münch (2018) maintain an individual gaussian belief state of catch expectations that each boat updates as it explores and lands fish (blurring the line between discrete-choice and agent-based model).

While some agent-based models also make perfect-knowledge assumptions (e.g, Gao & Hailu (2011), Elliston & Cao (2006)), most simulate information and exploration explicitly. Little & McDonald (2007) and Dorn (2001) use Kalman filters as a metaphor for fishermen's inference and learning. Fishers in Dorn (2001) use fixed thresholds to decide whether to fish or search while Little and McDonald (2007) agents always exploit the most promising area (given the Kalman filters' predictions). Bastardie et al. (2013)

expands on the threshold-based approach to multiple dimensions by using decision trees.

Such a large diversity in decision-making theories raises two questions: first, do any these complications matter at all? Second, how should we select the “best” among them? Coding them side by side in POSEIDON allows us to answer both.

We describe the fishery and available data in section 2. We enumerate the components of our model in section 3. In section 4 we describe how we calibrate and validate the model and in section 5 we describe the main results. We describe the sensitivity analysis in section 6. Finally, we discuss our main findings, assumptions and caveats in section 7 and conclude in section 8.

2. The fishery

2.1 The DTS fishery

The West Coast groundfish fishery is a multispecies fishery that lands over 90 species and has undergone significant management changes over the past 15 years. The fishery operates from southernmost California to northernmost Washington in coastal waters mostly within 100km of the shore. Species caught in the West Coast groundfish fishery are demersal (living near the bottom), including flatfish (e.g., sole), roundfish

(e.g., sablefish, *Anoplopoma fimbria*), and rockfish (*Sebastes sp.*). Primary target species are petrale (*Eopsetta jordani*) and Dover sole (*Solea solea*), shortspine (*Sebastolobus alascanus*) and longspine (*Sebastolobus altivelis*) thornyheads, arrowtooth flounder (*Atheresthes stomias*), and Pacific hake (a.k.a. whiting, *Merluccius productus*) (Miller & Deacon, 2017), though we focus only on non-whiting groundfish activities.

The species composition of catch is strongly related to fishing grounds. Ocean depth and demersal habitat type (e.g., rock, mud, sand, gravel) are major drivers of the species composition. The vast majority of the catch in the west coast groundfish fishery comes from trawl gear, which is largely indiscriminate within its tow-path (Bellman & Heery, 2013). A few vessels also use ‘fixed gears’ (e.g., hook and line, pots and traps), which are more selective (Miller & Deacon, 2017), typically to target sablefish.

The West Coast groundfish fishery is ideal to validate our model because it is relatively data rich and has a diverse management history, which includes trip limits, spatial and temporal closures, gear restrictions, limits on entry, and most recently, individual transferable quotas (ITQs; the Pacific Fishery Management Council uses the terminology individual fishing quotas – IFQs- though quotas can be transferred Miller & Deacon, 2017).

Here, we model the dynamics of one component of the groundfish trawl fishery that targets the ‘DTS’ complex (dover sole-thornyhead-sablefish’). Dover sole, sablefish, longspine thornyhead and shortspine thornyhead made up over 25% of the annual

revenue from the groundfish fishery (Pacific States Marine Fisheries Commission, [2017](#)). We modelled the dynamics of these four species plus a bycatch species, yelloweye rockfish, because of its historically high level of depletion, overfishing, slow rebuilding, and role as a constraining species in the groundfish fishery (small amount of quota available for it).

Many rockfish species were harvested during the 1980s and 1990s at rates that we now believe were unsustainably high. Those excessive harvests culminated in the U.S. government declaring the West Coast groundfish fishery an economic disaster in 2000. In response, the Pacific Fishery Management Council implemented spatial management and gear restrictions. Rockfish conservation areas (RCAs) were established in 2002 to reduce bycatch of overfished species (particularly darkblotched [*Sebastes crameri*] and canary rockfish [*Sebastes pinniger*]). The RCA boundaries shift within and across years but are generally close to the shelf, roughly excluding depths of between 100-275m along the entire West Coast to trawl gear. The Council also implemented gear restrictions (on trawl footrope sizes) to protect rocky habitats shoreward of the RCAs and required selective flatfish gear that reduces the catch of rockfish in these areas (Bellman, Heppell, & Goldfinger, [2005](#)). In 2003, the U.S. government began a vessel buyback program to reduce fishing capacity (Miller & Deacon, [2017](#)). In 2011, the fishery transitioned to multispecies Individual Fishing Quotas (IFQs)-a catch-share program--leaving the RCAs in place. This policy change allocated the coast-wide catch limits for each species among individual vessel owners and required all vessels to have

independent observers on board to monitor bycatch, discards, and interactions with protected species.

2.2 Data available

The DTS fishery is data rich. We use biological data from stock assessments and economic data describing fleet characteristics to parametrize the model while we use data on fishing locations and outcome to calibrate and validate it.

We parametrize the biological layer in section 3.2. Its two main sources are the stock assessment models developed by NOAA scientists at the NW Fisheries Science Center for the Pacific Fishery Management Council (Hicks & Wetzel, [2011](#); Stephens & Taylor, [2013](#); Stewart, Thorson, & Wetzel, [2011](#); Taylor & Stephens, [2013](#); Taylor & Wetzel, [2011](#)) and habitat suitability maps of the California Current produced as part of the Essential Fish Habitat synthesis (National Marine Fisheries Service, [2013](#)).

We parametrize the fleet in section 3.3. Our main data source is the catcher vessel report (Steiner et al., [2017](#)) which presents economic information for every ITQ participant, catcher-processor, catcher vessel, mothership, first receiver, and shore-based processor. We also include additional data from the catcher vessel report website (Northwest Fisheries Science Center, [2017](#)) and bespoke specialized data requests.

We describe the fishing outcomes we calibrated against in section 4.1.2. These combine quota attainment rate and landing observations (National Marine Fisheries

Service, [2017](#); Somers, Lee, Jannot, & J., [2016](#)), but include also more observations from the catcher vessel report and aggregate trip information from the Pacific Fisheries Information Network (PacFIN) logbook data from California and Oregon. We calibrated fishing decisions against logbook data for all boats landing in California and Oregon ports between 2011 and 2014 as shown in section 4.1.1.

3. Model

3.1 Geography

The model runs over a grid 50 cells in the east-west direction and 120 cells in the north-south direction. Each cell is a square of size 15.16km. The top left cell of the map is off Washington coast on the border with Canada (lat: 48.194, long: -125.878), the bottom right cell is on the Californian sea border with Mexico (lat: 32.017, long: -117.403). Distances between cells are Cartesian over their UTM 10 coordinates.

The map is also more coarsely divided into statistical areas: one-degree latitude and one degree-longitude rectangles. These areas have no biological significance and are in fact just groups of map cells but they matter for some decision-making algorithms who generalize space coarsely (see section 3.4) and the logbook analysis carried out in section 4.1.

3.2 Biology

The model tracks the number of fish per cell, by species, age and sex. Each cell of the map (excluding land ones) contains a set of vectors $\{A_{1,male}, A_{1,female}, A_{2,male}, \dots, A_{5,female}\}$ (subscript denotes species,sex) where:

$$A_{i,male} = \begin{bmatrix} \text{No.of male fish of species } i \text{ of age } 0 & \text{No.of male fish of species } i \text{ of age } 1 & \dots & \text{No.of male fish of species } i \text{ of age } 5 \end{bmatrix}$$

Coast-wide abundance comes from stock assessment models. We populate each map cell by distributing coast-wide abundance in proportion to the relative cell abundance predicted by the Essential Fish Habitat synthesis effort (National Marine Fisheries Service, 2013). These abundance maps were generated by fitting spatio-temporal models to the NWFSC groundfish trawl survey and habitat data from 2003-2011 (Shelton, Thorson, Ward, & Feist, 2014).

We assume that all fish in the same age and sex bin are of the same weight and length.

We compute length at age L_a by

$$L_a = L_\infty (1 - e^{-k(a-a_0)})$$

We convert length to weight by:

$$W = \alpha_w L^{\beta_w}$$

Where α_w and β_w are allometric weight parameters, L_∞ and k are the Von Bertalanffy growth parameters and these differ for each species and each sex.

Through the weight vector W we can compute the spawning stock biomass (SS) for each species as:

$$SS_i = W_{i,female}^T (M_i \circ A_{i,female})$$

Where M_i is the vector of maturity for each age bin and \circ is the Hadamard (element-wise) product. We compute maturity as a function of length l as:

$$M(l) = \frac{1}{1 + e^{-\gamma(l-L_{50})}}$$

Where L_{50} is the inflection point and γ is the slope of the maturity curve and differ for each species.

Four events modify abundance within a cell: recruitment, ageing, natural mortality and fishing mortality. Fishing mortality is driven by the fleet component of the model.

Recruitment, age and natural mortality are global, deterministic and yearly events.

Mathematically we update each element of each abundance vector as follows:

No. fish of species i of age a = $\{R \text{ if } a = 0 \ e^{-z} (\text{No. fish of species } i \text{ of age } a - 1) \text{ if } 0 < a \leq \text{maxage} \ 0 \text{ if } a > \text{maxage}\}$
 Here, R is the number of recruits each year and M is the annual natural mortality rate.

Recruitment follows the Beverton-Holt form, parameterized in terms of steepness (h), unfished recruitment (R_0), recruitment potential (ϕ) and spawning stock biomass (SS)

$$R = \frac{4hR_0SS}{R_0\phi(1-h) + (5h-1)SS}$$

We aggregate each cell's spawning stock biomass (SS) into a single coast-wide stock before computing the number of yearly recruits. We assume half of the recruits are male. We distribute recruits in each cell proportionally to the original allocation of abundance at the beginning of the simulation.

We did not model ecological interactions between fish (e.g., predator-prey, resource competition). There is no fish movement but we test the model sensitivity to it in section 4.4 of the appendix.

Sablefish is targeted by other fisheries on the US West Coast. To simulate this, at the end of each year each cell is targeted at random until the total non DTS catch reported for that year has been removed (total catch as reported in the PacFIN landings database). This allows us to integrate exogenous fishing mortality without explicitly modelling the boats that caused it.

3.3 Fleet

We simulate the fishing behaviour of all 90 trawlers of the West Coast DTS fishery. Vessels were distributed by port according to the 2010 census available from the Northwest Fisheries Science Center's catcher vessel report (Steiner et al., 2017). The northern-most port is Seattle, Washington; the southern-most port is Monterey, California. In order to better understand the dynamics of various behaviour models the fleet was homogenised with respect to quota allocation, hold size, and fishing ability. Each fisher is given the same proportion of quota, while common hold size and fishing ability are calibrated in section 4.

A boat fishing for one hour, catches a set of vectors x_1, \dots, x_n determined by:

$$x_i = \begin{bmatrix} q_i (s_i \circ r_i \circ A_i) \\ \text{No.of fish caught of species } i \text{ of age } 0 \\ \text{No.of fish caught of species } i \text{ of age } 1 \\ \vdots \\ \text{No.of fish caught of species } i \text{ of age } n \end{bmatrix}$$

Where q_i is the catchability scalar for species i , s_i is the selectivity vector, r_i is the

retention vector, A_i is the number of fish per age bin in the cell being fished and \circ is

the element-wise product. The selectivity (s_i) and retention (r_i) vectors are fixed to the

stock assessment values; the abundance vector (A_i) is a simulation's state variable and

the catchability scalar q_i is a free parameter. We define catchability differently from the stock assessments: there it represents the probability of catching each fish in the ocean, here it is the probability of catching only the fish within a trawled cell.

Simulated boats sell fish by weight and the price is fixed to the average prices observed between 2011-2014. We assume perfectly elastic demand and sale prices are constant between ports. Table 1 shows the ex-vessel price per kg landed. Notice that sablefish price does not include fixed gear premiums that are usually paid to other west coast boats targeting sablefish that are not part of the DTS fishery. Besides the five species we model, trawlers in the DTS fishery often catch many other species in small amounts. The catcher-vessel data reports that 32.21% of the catch (by weight) in the DTS fishery is made of other species not present in the model; we simulate this by assuming that same proportion of catches made within the simulation is "miscellaneous" catch which is sold at the observed 2011-2014 prices.

Table 1. The ex-vessel price (that is the price paid to a boat when landing fish) in the model.

Species Landed	Ex Vessel Price (\$/kg)
Dover Sole	0.67
Sablefish	4.32
Shortspine	1.04

Longspine	1.04
Yelloweye	1.08
Miscellaneous	1.76

Lian, Singh, & Weninger (2009) computed the physical effort limits for groundfish trawlers on the west coast to be 170 days at sea. This implicitly accounts both for repairs, weather conditions and fishing seasonality. Most DTS fishermen also participate in the shrimp fishery for an average of 50 days a year. Following this, each boat can be at sea for a maximum of 120 days. Each trip, fishers trawl in a target cell until either their hold is full or six days have passed. This maximum is only binding when agents (vessels) are acting completely at random (as in the “random” algorithm described below). All other agents perform shorter trips.

Each boat steams at 16 kph consuming 3.54 L/km of gas; each hour spent trawling consumes 57L of gas (using Toft et al. (2011) estimates) and all trawling happens within the same cell. Gas price changes yearly. Within a year, gas price is fixed to the average yearly price observed in California for that year. We assume gas costs are the same in each port. On top of fuel expenditures, each hour at sea costs a boat \$165 (including crew and captain fees)¹.

¹ This was obtained from economic data collected by the Northwest Fisheries Science Center and publically available in the FISHEyE database (<https://dataexplorer.northwestscience.fisheries.noaa.gov/fisheye/>). We used the 4 year-average median variable cost, subtracted the average fuel costs and divided by the average days at sea.

Fishers are subject to two sets of regulations from 2011: Rockfish Conservation Areas and individual tradable quotas. Rockfish Conservation Areas are areas of the sea around the coast where fishing is not allowed. We simplify their boundaries by assuming they cover every cell of the map with depth of 275m or less.

In the model, each year, each fisher is allocated a set of individual quotas (referred to as quotas throughout this document). The quotas represent the maximum weight of each species that a fisher can land that year, derived as a proportion of the total sustainable catch for the entire fishery. Fishers can lease quotas to one another.

Fishers that run out of one quota and do not lease more are not allowed at sea until the next year. The quota leasing market functions as an order book: each day it reveals bid and ask prices for each quota from each fisher and matches crossing orders until every feasible trade occurs (always at the ask price with minimum price of 0.05\$). Quotas cannot be sold permanently in the model as a moratorium against it was in place until 2014. There is a single global order book and quotas can be leased between fishers of different ports with no transaction cost or friction.

Expectations and reservation prices are as derived in the original appendix of POSEIDON: λ_i the reservation price for quota of species i , given expected daily catches c_i , profit per catch Π_i and current price of quotas p_i is given by:

$$\lambda_i = (\Pi_i + \sum_{j \neq i} \frac{c_j}{c_i} (\Pi_j - p_j)) P$$

Where P is the probability of having to use the quota, and is equal to the probability of c_i being above total quota currently held by the fisher divided by the numbers of days left in the season. The distribution of c_i is assumed normal with mean and standard deviation equal to the observed moving average and standard deviation for the past 365 days.

The aggregate quota available is fixed at the average 2011-2014 observations summarised in Table 2. We removed from sablefish quotas those that do not belong to the DTS fishery: quotas south of the 36N parallel as well as the landings achieved by the separate fixed gear fishery. We test this assumption during sensitivity analysis in the appendix.

Table 2. Yearly quotas to be shared among fishers each year.

Species	Yearly Quota (mt)
Dover Sole	22,234.5
Sablefish	1,606 ²
Shortspine	1,4816
Longspine	1,966.25
Yelloweye	0.6

² We Unifying north and south quota allocations and not removing fixed gear landings would generate 2,725t of sablefish available. Which is what is allocated in the sensitivity test.

A boat may decide not fish for the entire season if it manages to lease all its quotas or if it quits the fishery permanently. A boat that makes two years of consecutive losses will quit the fishery permanently. Errend et al (2017) show that on average for 10% of the vessels operating costs are higher than revenues, but no information is provided on how many boats incur losses consistently across the years. Even after quitting permanently boats still own quotas and lease them in the ITQ market to active fishers.

All model runs have one initialization year: agents fish for a year following all the proper regulations, after which the biology layer is reset and the real model run starts. This helps adaptive and RUM agents to initialize their memory as well as allow agents to form expectations about daily catches which are needed to price quotas.

3.4 Decision-making

Fishers need to decide which cell to trawl each trip. Here we implement a set of alternative decision making algorithms and let the available data compare the appropriateness of each. The algorithms vary in terms of rationality, information available, parameters needed and how they are calibrated. We group them in 3 sets: adaptive agents, statistical agents and standard assumption agents. Table 3 describes the algorithms and how they were calibrated (see section 4.2), table 4 describes the information they observe and process, table 5 describes the parameters of the adaptive agents as well as their function.

Table 3: A brief description of each algorithm used for decision making, split into three categories: adaptive, statistical and standard assumptions

Algorithm	Description	# of parameters	Calibration Target
<i>Adaptive</i>			
Social Annealing	Always fish the same cell unless they are making less than k% of the average fishery's profits, at which point they explore	1	Logbook error
Heatmap	Progressively build a statistical heatmap as they experience where profits are higher, and target its peaks	4	Logbook error
EEl (explore-exploit-imitate)	Have a fixed probability of exploring a neighbouring cell or copying the location of other fishers who are making more profits; the "uncalibrated" version has its parameters set to the POSEIDON default of 20% exploration chance, 5 map cell radius and 100% imitation rate.	3	Logbook error
Bandit (epsilon-greedy)	Fixed probability of exploring a random new statistical area, otherwise fish where the average profits observed are the highest	2	Logbook error
<i>Standard Assumptions</i>			
Random	Every trip, pick cell at random	-	-
Perfect – Cell	Always choose the cell that will generate most profits	-	-
Perfect – Statistical Area	Chooses the statistical area that generates the highest profit but does so probabilistically using a logit function.	-	-
<i>Statistical</i>			
RUM	Uses discrete-choice model choosing statistical areas from habit, distance, revenue and cpue (sole, yelloweye and sablefish) variables	7	Outcome error
RUM Fleetwide	As above, but with cpue and revenue information shared across the fleet	7	Outcome error
RUM Precise	As RUM, but choose POSEIDON cells (more precise)	7	Outcome error
Historical	Fishes in statistical areas in proportion to how often they were fished in the logbook data	-	-
Logit	Uses discrete-choice model choosing statistical areas from habit, distance and intercepts variables from logbook data	-	-

Table 4: a description of the parameters of each of the adaptive algorithms

Algorithm	Parameter	What it does
Bandit	α	Exponential moving average parameter; weighs new observation with current expected average profits
	ε	Probability of exploring new statistical area each trip
EEl	ε	Probability of exploring new cell each trip
	l	Probability of copying the location of somebody making more profits for each trip (when not exploring and there is somebody making more)
	δ	Von Neumann neighbourhood size in terms of cells the fisher will pick his next trip destination from when exploring
Heatmap	ε	Probability of exploring new cell each trip
	δ	Von Neumann neighbourhood size in terms of cells the fisher will pick his next trip destination from when exploring
	α	Forgetting factor of the Kernel regression building the heatmap; discounting, each time new information is produced, older observations
	Bandwidth	Kernel regression parameter describing how far in the distance should an observation generalize (for example, how much should an observation in cell 1,1 update our beliefs about profit in cell 3,3)

Simulated Annealing	k	% of profits made compared to fishery's average above which the agent stops exploring
	δ	Von Neumann neighbourhood size in terms of cells the fisher will pick his next trip destination from when exploring

Table 5: a brief description of the information available and used by each decision-making algorithm

Algorithm	Information Available
Adaptive Models	
Social Annealing	Only knows average profits made within the entire DTS fishery
Heatmap	Knows profits made by the fisher last trip as well as the profits and locations of trips made by 2 other fishers randomly chosen from the same port. Feeds observed profits to update a kernel regression representing the agent's belief about future profitability'
EEl (explore-exploit-imitate)	Only knows the last profits made as well as the profits and locations made by 2 other random fishers from the same port. Has no memory except for the latest trip made..
Bandit (epsilon-greedy)	Observes only own profits and choices, keeps track of profits made in each statistical area using exponential moving averages
Standard Assumption Models	
Random	Knows nothing
Perfect – Cell	Knows the profitability of each cell of the map perfectly but does not take into account the possible future actions of other fishers
Perfect – Statistical Area	Knows the profitability of one random cell in each statistical area for each trip made; does not take into account future actions by other fishers
Statistical Models	
RUM	Observes only own revenues, catches and profits. Keeps track of average CPUE, revenue and times the fisher has visited each statistical area in the past 365 simulated days.
RUM Fleetwide	As above, but memory is shared fleet-wide
RUM Precise	As RUM, but keeps track of information at cell-level rather than statistical area.
Historical	Knows real distribution of fishing for each statistical area and endeavors to reproduce it
Logit	Knows distance and keeps track of # of time the agent fished each statistical area in the past 365 simulated days

We do not presume that any of the algorithms perfectly replicate the decision-making process of actual fishers. Instead, our goal is to compare the performance of these algorithms in terms of how well they capture individual and aggregate behaviour and outcomes observed in the DTS fleet.

Adaptive agents are introduced in a conceptual format in Carrella et al.(2018) and we show their pseudocode in the appendix. For EEI agents we also test the default setting (uncalibrated) version of the decision making algorithm to check its robustness to misspecification.

Statistical agents are either fit to data beforehand (the *historical* and *logit* agents) or use a multinomial regression structure³ (*RUM*, *RUM fleetwide* and *RUM precise* agents) whose parameters minimize the outcome error defined in section 4.

Standard assumption agents are either perfect knowledge or purely random. *Random* agents are useful to prove that behaviour matters in this model. *Perfect* agents know the location of all the fish and how profitable fishing will be in each area before setting off.

4. Calibration and validation

Calibration means model fitting: changing parameters to minimise the distance between model output and real data (Windrum, Fagiolo, & Moneta, 2007 is a review of this problem for agent-based models). Validation means testing the quality of out-of-sample predictions (Schulze et al., 2017 calls this "output corroboration", but as noted in Augusiak, Van den Brink, & Grimm, 2014, it is a procedure that goes by many names).

³ That is, the probability of choosing location i to fish next trip is the softmax function where X is distance, revenue, CPUE for 3 species, habit and an intercept (which acts as an exploration gauge) and the β vector is constant for all alternatives

For both procedures we first generate a set of "summary statistics" that describe the real data (Hartig, Calabrese, Reineking, Wiegand, & Huth, [2011](#)). We then simulate the same set of summary statistics from the model. The closer the simulated statistics are to the real ones, the lower the error.

4. 1 Summary Statistics

4.1.1 Fishing Decisions

Logbook data contains the length, duration, catch (by species) and location of each trip made by each boat from California and Oregon, self-reported to the Pacific Fishery Information Network (PacFIN) by the vessel captains. We condense this large set of observations into a few key metrics. The fewer and more informative these metrics are, the easier our calibration. We obtain these metrics by fitting a discrete-choice model to the logbook dataset.

Discrete-choice models are often used in fisheries management to study why and how captains choose where to fish (Abbott & Wilen, [2011](#); Holland & Sutinen, [2000](#); Hunt, [2005](#); Hunt, Boots, & Boxall, [2007](#)). We first impose a grid of statistical areas over the California Current, each one-degree latitude by one-degree longitude as shown in figure 1. These grids are larger than those used with the ABM simulations which requires us to aggregate samples over a larger area in order to estimate the discrete-choice model. We then fit a multinomial logistic regression that predicts which statistical area each vessel goes to next given its history.




Figure 1 Map of defined statistical areas off the West Coast of the United States, which represent discrete fishing location choices in the revealed preference panel dataset, and in the output of the POSEIDON model application to the West Coast groundfish fishery.

We fit our logistic regression using only 2 variables (excluding intercepts): habit (times each area was visited in the past year) and distance from port. We choose this model for three reasons. First, the regression has an in-sample success rate of 72% (i.e. probability of predicting correctly which of the 32 statistical areas the fisher will go next). Second, the very high habit coefficient summarises a key dynamic of the fishery: agents often return to the same statistical area. Third, parsimoniously, we sought to fit a statistical model with as few parameters as possible for ease of calibration.

Table 6 shows the coefficients of the logistic regression. All intercepts except one are insignificant (and the significant effect size is small compared to habit and distance). The farther an area from port, the less likely it will be chosen, while areas that have been visited before are more likely to be chosen again. We use these parameters to calibrate our agents in section 4.2.

Table 6 Multinomial site choice model β for all trips from 2012 (trips in 2011 used to initialize the habit parameter). Distance is the kilometres between port and centre of statistical area. Habit is number of trips in the same area in the past 365 days. Bolded are the coefficients that are significant.

Parameter	Estimate	Std. Error	t-value	Pr(> t)
Distance	-0.021	0.001	-15.094	<0.001
Habit	0.213	0.013	16.261	<0.001
51:(intercept)	-0.029	0.639	-0.045	0.964
61:(intercept)	1.983	0.576	3.443	0.001
62:(intercept)	0.953	0.529	1.801	0.072
63:(intercept)	0.509	0.476	1.071	0.284
73:(intercept)	-16.018	5752.487	-0.003	0.998
85:(intercept)	0.825	3.966	0.208	0.835
86:(intercept)	-16.758	5682.723	-0.003	0.998
97:(intercept)	2.226	3.964	0.561	0.574
98:(intercept)	-0.138	3.997	-0.035	0.972
109:(intercept)	-16.328	6539.433	-0.002	0.998
110:(intercept)	1.869	3.984	0.469	0.639
122:(intercept)	1.802	4.103	0.439	0.66
134:(intercept)	2.693	4.121	0.653	0.513

135:(intercept)	1.959	4.107	0.477	0.633
147:(intercept)	3.035	4.117	0.737	0.461
148:(intercept)	2.434	4.116	0.591	0.554
160:(intercept)	2.842	4.135	0.687	0.492
161:(intercept)	-0.115	4.246	-0.027	0.978
173:(intercept)	3.614	4.143	0.872	0.383
174:(intercept)	1.706	4.152	0.411	0.681
186:(intercept)	3.105	4.147	0.749	0.454
187:(intercept)	1.335	4.155	0.321	0.748
198:(intercept)	-17.351	7252.268	-0.002	0.998
199:(intercept)	2.59	4.148	0.624	0.532
200:(intercept)	1.068	4.149	0.258	0.797
212:(intercept)	1.497	4.151	0.361	0.718
213:(intercept)	-0.74	4.163	-0.178	0.859
225:(intercept)	2.388	4.154	0.575	0.565
226:(intercept)	-18.703	6210.947	-0.003	0.998
237:(intercept)	-15.193	6316.336	-0.002	0.998
238:(intercept)	3.905	4.162	0.938	0.348

4.1.2 Fishing outcomes

Table 7 lists the aggregate observations about the fishery we want our model to reproduce. Notice that while the empirical mean cited refers to the period 2011-2014, the standard deviation refers to the between years standard deviations (as we obviously have only one possible observation of the 2011-2014 mean).

Table 7 The aggregate observations we would like our model to be able to predict. All means refer to the period 2011-2014. All standard deviations are the empirical deviations of the yearly mean observed, except for the data derived from the logbook data where we are using the standard deviation of the whole sample.

Measure	Empirical Mean	Empirical Standard Deviation	Source
Boat Yearly Profits (\$)	118,552\$	21,331.00	Steiner et al. (2017)
Yearly Hours Out	999.936	120	Steiner et al. (2017)
Sole Quota Attainment	33.25%	3.09%	Somers et al. (2017)
Sablefish Quota Attainment	83.65%	6.18%	Somers et al. (2017)
Shortspine Quota Attainment	52.50%	5.00%	Somers et al. (2017)
Longspine Quota Attainment	51.50%	5.00%	Somers et al. (2017)
Yelloweye Quota Attainment	7.00%	2.00%	Somers et al. (2017)

Trip Duration (hr)	69.1	33	(Pacific States Marine Fisheries Commission, 2017)
Distance port to 1st trawl (km)	90.89	32	(Pacific States Marine Fisheries Commission, 2017)

4.2 How to calibrate

We selected a few key summary statistics from data. The model simulates new summary statistics. Error is the distance between real and simulated summary statistics. Calibration involves aggregating these errors into a single number and tuning the model parameters to minimize it. We define two such measures: logbook error and outcome error.

Logbook error measures the difference between the real logit regression parameters β from table 4 and the coefficients $\hat{\beta}_i$ produced by tracing simulated agents within POSEIDON, collecting their logbooks and applying the same discrete-choice regression to this simulated logbook. Mathematically this is:

$$\text{Logbook Error} = \sum_i \frac{|\beta_i - \hat{\beta}_i|}{\sigma_i}$$

Where σ_i is the standard error of each coefficient. The discrete-choice model fit contains area-specific intercepts as in table 6 but their differences are not part of the logbook error⁴.

⁴ This is to avoid the problem of agents visiting statistical areas where no real trip was observed and for which no coefficient and standard deviation is defined.

This is an example of indirect inference (Gourieroux, Monfort, & Renault, 1993). Indirect inference has been applied to agent-based models in the past, particularly for financial and economic simulations (Richiardi et al., 2006; Zhao, 2010). An idiosyncratic version of indirect inference is present in the fishery model by Cenek & Franklin (2017) where the authors use linear regressions as auxiliary models.

To our knowledge indirect inference by discrete-choice model has never been done in spite of this being a very straightforward solution to the issue of comparing simulated and empirical maps (see section 6 of Stow et al., 2009 where the β of our auxiliary model can be understood as just another higher order spatial feature).

The outcome error is the difference between the observed fishing outcomes summarised in table 5 and the fishing outcomes generated by the agent-based model. More precisely the error is defined as:

$$Outcome\ Error = \sum_i \frac{|y_i - \hat{y}_i|}{\sigma_i}$$

where y_i is the empirical mean of a fishing outcome, σ_i is its standard deviation and \hat{y}_i is a simulated fishing outcome.

The weakness of this approach is that it implicitly weighs some features more than others. Because landings data have lower variation coefficients than profits, we penalize prediction errors in landings more than in profits. Moreover, because our landings are defined in terms of quota attainment and percentage points, we value accuracy in yelloweye attainment (for which less than a ton of quota is available each year) just as much as sablefish accuracy (the main target species). This is however an unavoidable

consequence of aggregating multiple patterns into a single error number (Badham, Jansen, Shardlow, & French, 2017).

The model has 6 non-behavioural parameters: a vector of catchabilities

$Q = (q_{sablefish} \ q_{sole} \ q_{shortspine} \ q_{longspine} \ q_{yelloweye})$ and the maximum hold H size of each boat. Some decision-making algorithms have additional parameters to tune (listed in section 3.4).

We calibrate our model in two steps. First, using *historical* agents (agents that probabilistically go where boats empirically fished), we calibrate the catchability vector Q and the hold size H by minimizing the outcome error. Second, given the catchability and hold size from step one, we calibrate each decision-making algorithm separately by minimizing the logbook error. The exception to step 2 are the *RUM*, *RUM Fleetwide*, *RUM Precise* algorithms who minimize outcome error in step 2 as well, since minimizing logbook error would make them equal to the *logit* agents⁵.

We minimize the errors in each step by using a Bayesian optimiser (Shahriari, Swersky, Wang, Adams, & De Freitas, 2016; Snoek, Larochelle, & Adams, 2012). In essence this is an evolution of the BACCO (Bayesian analysis of computer code output) techniques that are popular in environmental agent-based models (O'Hagan, 2006; Parry, Topping, Kennedy, Boatman, & Murray, 2013; Uusitalo, Lehtikoinen, Helle, & Myrberg, 2015).

⁵ This is because RUM agents are the same multinomial logistic regression formulas as the *Logit* agents but with more parameters.

An optimiser calibrates the model efficiently but stresses the implicit weighing due to error aggregation. That is, the optimiser trades large increases in error for variables with high σ in exchange for minor improvements in variables with lower standard deviations.

4.3 How to validate

We have outcome summary statistics for 2015 and 2016 which we did not use for calibration. We simulate those two years with each calibrated model and compare simulated summary statistics against real ones.

First, we look for “anti-patterns”: decision-making algorithms that generate consistent behaviour at odds with the outcomes and summary statistics we observe (as for example a decision-making algorithm that consistently cause the collapse of the fishery). We reject algorithms that produce “anti-patterns”.

Second, we rank algorithms by the 2015-2016 out of sample outcome error (which we call here validation error). As in the calibration section, validation error is an aggregate of each summary statistic error weighted by their observed standard deviation. As with calibration, validation error is sensitive to the way weights are specified (in particular regarding yelloweye rockfish errors where total quota is very limited). To avoid depending on a single set of weights, we also employ three automatic model selection schemes, each producing a different weighing structure: approximate Bayesian computation(ABC) by rejection (Csilléry, François & Blum, 2012), ABC by random

forests (Pudle et al, 2015) and elastic nets (Carrella, Bailey & Madsen, 2018). We then compare rankings produced by all methods.

Rejection ABC normalizes all outcome errors by their standard deviation across all simulations and selects only the 10% of the simulations with the smallest normalized validation error. It then assigns the probability of each decision-making algorithm being correct as the proportion of its simulations that were among the top 10%.

ABC by random forest uses a random forest classifier (Breiman, 2001) to weigh each summary statistic error in a way that makes it easier to discriminate between decision-making algorithms. Intuitively the approach is not to look for the decision making algorithm that minimizes a validation error, rather it is to discover what patterns in each simulated summary statistic is indicative of the decision making algorithm that generated it and then use this knowledge to decide which decision making algorithm is more likely to have generated the real summary statistics.

Random forests are non-linear and non-parametric; a linear, regularized, parametric alternative is elastic-nets (Zou and Hastie, 2005). Again this turns the problem of finding the weights for the aggregate validation error into a classification problem of finding the weights that better discriminate among decision-making algorithms.

We know of no guidelines on how many simulations to run in order to make model selection is effective. We therefore follow a standard power analysis criterion (Lipsey, 1990; Seri & Secchi, 2017): targeting a statistical power of 95%, effect size of 0.2 and α

of 0.05 Bonferroni's corrected for 11 comparisons (comparing the top strategy against all the others) we run the model 1006 times for each decision-making algorithm.

We also validate the model in two other ways. First, in section 2.2 of the appendix, we look at qualitative predictions of other fishing outcomes we didn't explicitly calibrate against. Second, in section 3 of the appendix, we look at behavioural rankings after modifying some assumptions of the model (changing the starting date to 2007, changing the distribution of fish, ignore one-trip boats when computing profits and unifying the sablefish market).

Note that the validation years 2015-2016 were not radically different in outcomes from the calibration years that preceded them. This weakens the power of any test to be able to distinguish good decision algorithms from bad ones. Ideally we would use as validation the aftermath of a large policy shock (as suggested in Reimer et al, 2017). We unfortunately do not have a such "nonrandom holdout sample" (Keane and Wolpin, 2007).

5. Results




Figure 2: *Box-plot showing the logbook and outcome errors for 1006 runs of each decision-making algorithm. Red shaded algorithms are statistical, blue shaded algorithms are adaptive and green shaded algorithms represent traditional assumptions. Adaptive (blue) algorithms were trained to minimize logbook error (left), statistical algorithms (red) were trained to minimize outcome error; traditional agents were not trained. Random boats quit before a logbook regression can be run, RUM precise and Perfect Cell stick to fishing in a single statistical area and their logit regression is unidentifiable.*




Figure 3: *Box-plot showing the validation outcome errors for 1006 runs of each decision-making algorithm for 2 years of withheld data. Red shaded algorithms are statistical, blue shaded algorithms are adaptive and green shaded algorithms represent traditional assumptions.*

For each decision making algorithm, figure 2 shows the calibration (in-sample) errors for all algorithms, figure 3 shows validation (out-of-sample) errors. Table 8 lists the weighted errors for each individual statistic and section 2.1 of the appendix contains the plots of the produced summary statistic for each algorithm.

Table 8: average error for each summary statistic and each algorithm for the 2015-2016 out of sample period.

Algorithm	Profit	Hours out	Sole	Sablefish	Long Thornyheads	Short thornyheads	Rockfish	Validation Error
Logit	3.31	1.44	5.48	0.95	0.39	3.82	3.52	18.87
Historical	2.93	0.35	3.52	0.8	0.21	3	3.52	14.42
RUM - precise	4.24	0.78	5.76	0.77	2.42	5.25	3.52	22.71
RUM	2.36	0.88	3.75	0.24	0.6	3.42	3.52	14.97

RUM Fleetwide	3.25	2.86	5.04	1.51	0.39	4.27	3.52	21.08
Perfect - SA	0.38	1.61	2.49	0.28	1.32	3.67	34.92	44.57
Perfect - Cell	1.01	1.88	5.21	0.68	1.53	1.73	41.02	52.34
Random	6.3	6.64	9.76	14.02	7.17	9.78	3.52	57.19
Bandit	3.13	1.51	4.49	1.54	0.63	4.04	12.74	28.43
EEI	0.8	0.89	1.03	0.38	1.23	2.11	3.52	10.37
EEI - Uncalibrated	0.36	1.79	2.59	0.75	0.79	1.77	3.52	11.8
Heatmap	1.7	0.46	2.03	0.17	1.23	2.03	3.52	11.19
Social Annealing	5.75	6.88	6.92	8.73	5.38	7.71	3.52	45.51

Four algorithms, *random*, perfect at statistical area (*Perfect SA*), perfect at POSEIDON cell (*Perfect Cell*) and *social annealing*, generate “anti-patterns”: consistent outcomes we do not see in the real world.

Random agents fail to achieve any profits and cause the fishery to close down. This should be an obvious result but often in tightly calibrated environmental agent-based models the biological parameters drive all the results, drowning out human behaviour (Janssen, 2009; Schulze, Müller, Groeneveld, & Grimm, 2017). The failure of random agents shows that calibrating catchability is not enough to simulate the fishery; behaviour matters too.

Perfect SA agents achieves the highest profits for algorithms that choose statistical areas, *Perfect Cell* agents achieves the highest profits overall (see figure 4).

Both achieve some of these profits by purposefully targeting yelloweye rockfish either at the very beginning or the very end of the fishing season (see figure 5). They catch exactly the right amount and never fish out too much of the quota available. This behaviour does not occur in the real fishery (Holland, 2016).

Perfect Cell agents also fish areas where they can extract more dover sole catches for each unit of sablefish they land. Figure 6 shows quota attainment (% of quota landed over quota available) for Dover sole for each decision making algorithm. *Perfect Cell* agents catch 7 kg of Dover sole for each kg of sablefish, while *Historical* and *EEI* catch 3.24kg and 4.93kg, respectively.




Figure 4: The average profits made in the out-of-sample 2015-2016 period for each decision making algorithm. Red dashed line is the real observed 2015-2016 average.




Figure 5: The average percentage of yelloweye quota attained (quota landed divided by quota available) for each decision making algorithm in the out-of-sample 2015-2016 period. Red dashed line is the real observed 2015-2016 average. Partially due to its

rarity in most simulations no yelloweye rockfish is landed by 2015-2016.




Figure 6: The average percentage of Dover sole quota attained (quota landed divided by quota available) for each decision making algorithm in the out-of-sample 2015-2016 period

Social annealing agents generate a steady decline in catches, profits and active fishers. This is due to its behavioural parameters instantiating an easily satisfied agent with a small exploration range whose fishing efficiency declines over time as a result. We show in section 5.2 of the appendix that it is possible to find parameters that achieve low validation error for *social annealing* but these represent hyper competitive agents far from the satisficers (Simon, 1978) that the algorithm is meant to represent and test.

Other heuristics do not generate obvious anti-patterns and their ranking is therefore more subjective. However, both the validation error and the three model selection schemes described above select adaptive algorithms as more believable. Table 9 show the probabilities assigned by rejection ABC, random forest ABC and elastic-net classification. As expected no algorithm finds perfect or random behaviour likely.

Bandit	EEI	EEI Uncalibrated	Heatmap	Historical	Logit	Perfect - Cell	Perfect - SA	Random	RUM	RUM - precise	RUM Fleetwide	Social Annealing	Method
--------	-----	---------------------	---------	------------	-------	-------------------	-----------------	--------	-----	------------------	------------------	---------------------	--------

0.20 8	0.24 3	0.23 0	0.21 6	0.00 1	0.04 6	0.03 0	0	0	0.02 1	0.00 5	0	0	ABC
0	0.75 6	0.22 6	0.00 2	0	0	0.00 2	0	0	0.01 4	0	0	0	RF
0	0.99 5	0.00 1	0	0	0	0	0	0	0.00 4	0	0	0	Elastic nets

Table 9: the probabilities of each heuristic being “true” by each model selection method

RUM agents have low calibration but high validation error. This effect is due to the fishing efficiency of RUM agents decreasing over the simulation, becoming low by the 2015-2016 validation period.

In section 5.1 of the appendix we show that it is possible to have *RUM precise* agents with low validation error. To achieve this we need a strongly negative β coefficient for the habit variable. Indirectly then we are injecting into the RUM agent an exploration incentive to make it act like the other adaptive agents.

6. Sensitivity Analysis

Detailed sensitivity analysis results are in the appendix. Section 3 of the appendix contains four separate scenarios where some assumptions of the original model are modified. Section 4 of the appendix contains active non-linear tests (Miller, 1998) for both catchability, behavioural parameters and fish movement.

In the appendix section 3.1 we start the model in 2007 rather than 2011 to simulate the transition to ITQ. Adaptive algorithms still minimize validation error, with *heatmap* agents doing better than *EEI*.

In the appendix section 3.2 we modify the distribution of fish, using CPUE maps from the logbook data rather than the EFH dataset. In this sensitivity analysis adaptive agents do no better than statistical and standard assumption agents. All simulations however generate a new anti-pattern as shortspine becomes a constraining species for sablefish which results in quota prices 10 to 20 times above the observed values.

In the appendix section 3.3 we re-calibrate the model by counting only boats making two or more trips per year in the average profit summary statistic. We do so because boats that hold on to few quotas may still find one trip profitable but in general bias

down the average profits made by the fishery. In this scenario *heatmap* agents minimize validation error, however one of the three model selection algorithms (elastic nets) find the *historical* agents to be most accurate.

In the appendix section 3.4 we re-calibrate the model allowing both north and south sablefish quota to be available for the DTS fishery (ignoring therefore the contribution of fixed gear fishers). Ranking by both validation error and model selection algorithms reaffirm adaptive agents as more accurate, but validation error is in general inflated by too much sablefish being landed compared to reality.

The active non-linear tests for historical, EEI and heatmap agents find only marginal sensitivity to parameters. A search within the 10% interval for the optimal parameters find sensitivity of 11% for the historical agent, 6% for the EEI agent and 13% for the heatmap agent. Adding sablefish movement has no real effect on validation error.

7. Discussions

7.1 Simple adaptive rules describe fishing behaviour better

We showed that simple adaptive agents capture observed aggregate fleet-level patterns better than perfectly rational ones. The issue with perfect agents is one of proportions.

In Toft et al. (2011) and Kaplan et al. (2014) where the agents are port groups and geography is coarse, it is probably fine to assume agents are coarsely perfect. With higher resolution models where agents represent individual vessels, this assumption can generate “anti-patterns” (in our case high rates of yelloweye rockfish attainment, efficient Dover sole targeting and consistently higher profits).

The difference between perfect and adaptive agents is one of ability, not objectives.

Adaptive algorithms (with the exclusion of social annealers) are still profit maximizers.

They always prefer more profits and do not stop searching or fishing when an acceptability threshold is met. Adaptive agents do not represent a different attitude to risk than perfect agents either. By virtue of knowing everything, perfect agents do not ever incur risk. They are not risk-neutral or risk-prone, they are risk-free. Adaptive agents, instead, risk allocating the wrong amount of time exploring versus exploiting. This is simply an unknown cost in their profit function, and they are risk-neutral towards it.

7.2 Assumptions and caveats

We sought to build the simplest possible model for the US west coast groundfish fishery that could incorporate and approximate the data available to us. Keeping a model simple reduces free parameters and degrees of freedom, but targeting simplicity is just one possible heuristic (Edmonds & Moss, [2005](#)). The cost of this simplicity was that we did not model all of the complexity present in the West Coast groundfish fishery.

We chose to model geographic heterogeneity in the biological layer rather than focus on vessel heterogeneity, which may be just as important in the fishery. We did so because geographic heterogeneity is necessary for behaviour incentives to emerge. For a decision-making algorithm to perform well, variability must exist in the quality of fishing locations. Some areas must have higher abundance than others, and the risk of catching constraining species must also be unequal across areas. Only then can decision-making algorithms be tested, our primary objective for this study.

The comparison between adaptive agents and statistical ones is weak. We used a simple 2 parameter logit model to compare to the adaptive agents. We implemented more complicated RUM agents within the model but we calibrated them by minimizing outcome error within the model which proved sensitive to overfitting. Moreover this implementation of POSEIDON does not contain wind speed, currents and other variables that are common in discrete-choice models of fishers.

We did not model the entire value chain for groundfish (as Cooper & Jarre, [2017](#) do in their Hake fishery model, for example). In our model boats are price-takers facing constant prices. While the processing sector downstream is concentrated we did not model its monopoly power. Fishers in our model pick where and what to target without bargaining with their buyers. Modelling buyers may add geographical realism to the model as some ports may only be able to fulfill a limited number of orders.

Our model assumes fishers face a simplified cost structure. In reality crew and captains in the DTS fishery are in part paid a share of the revenue (see chapter 10 of Steiner et al., [2017](#)); buyback fees are paid as a percentage of landings while observers are paid in 24-hour increments. We subsumed all these into an hourly variable cost.

We assumed fixed catchability in our model; if agents want to change their catch composition they must pick a different cell to fish from. This assumption may be too restrictive, even within the limited time span of our model. Miller & Deacon ([2017](#)) show that some fishers may be fishing more at night or make shorter trawls to avoid bycatch.

We did not allow fishers to change home port. In our model agents may quit the fishery but are otherwise tied to their port. We made this assumption because every port in the DTS fishery has fewer boats in 2014 than in it had 2011. In the long term, however, it is possible that those who survive may move to more profitable areas.

While we assumed all agents were profit maximizers, Klein, Barbier, & Watson (2017) report that consistency, sustainability and neighbourliness may be just as important as profit in this fishery. The decision-making algorithms we presented can be adapted to maximize utility rather than profits, and future work could explore its implications.

We did not model alternative fisheries these boats may be employed in during the rest of the year: fixed gear, whiting or shrimps. We simply assumed agents have a maximum number of days to dedicate to the DTS fishery. This is a problem because there tends to be quota leftovers each simulated year, whose credibility depends on the profits boats can make in the alternative fisheries.

Events in the out of sample period we have chosen for validation (2015-2016) are not very different from those that preceded it. In replicating this approach to other areas we should look for cleaner breaks between calibration and validation periods, ideally due to sudden policy changes.

8. Conclusion

The push for more holistic approaches to management of natural resources requires new tools that include aspects of human and biophysical systems. To date, too few bioeconomic models have been used to inform decision making for fisheries problems, resulting in reactive, rather than proactive management responses. For fishery managers, tools that predict how individuals and fleets will respond to changes to regulations or harvested fish species would fill a critical gap.

POSEIDON is an agent-based model of fisheries that can begin to fill this gap. Here we calibrated the model, validated it and tested its sensitivity in a complex multispecies fishery on the west coast of the U.S. The DTS fishery is a relatively data rich example (stock assessments, fish habitat models, economic data, logbook data, complete observer coverage) and had a computationally tractable number of vessels. It provided an appropriate benchmark to test the performance of a suite of decision-making algorithms and sensitivities to assumptions about model structure.

This application fundamentally shows that behaviour matters in terms of outcomes even after keeping biological and fishing technology constant. This implies that agent-based models that implement multiple alternative behaviours can be useful to analyse the sensitivity of policies to the full range of possible adaptation techniques. More applications will follow which will help to generalize the rankings and sensitivities we identified in the DTS fishery.

Acknowledgments

This research is funded in part by the Oxford Martin School, the David and Lucile Packard Foundation, the Gordon and Betty Moore Foundation, the Walton Family Foundation, and Ocean Conservancy. We would like to thank the excellent reviewers at NOAA Fisheries (L. Pfeiffer, J. Hastie, M. McClure) for vastly improving the quality of this manuscript. This manuscript reflects the views of the authors and not NOAA Fisheries.

References

Abbott, J. K., & Wilen, J. E. (2011). Dissecting the tragedy: A spatial model of behavior in the commons. *Journal of Environmental Economics and Management*, 62(3), 386–401. <https://doi.org/10.1016/J.JEEM.2011.07.001>

Augusiak, J., Van den Brink, P. J., & Grimm, V. (2014). Merging validation and evaluation of ecological models to “evaluation”: A review of terminology and a practical approach. *Ecological Modelling*, 280, 117–128. <https://doi.org/10.1016/j.ecolmodel.2013.11.009>

Badham, J., Jansen, C., Shardlow, N., & French, T. (2017). Calibrating with Multiple Criteria: A Demonstration of Dominance. *Journal of Artificial Societies and Social Simulation*, 20(2), 11. <https://doi.org/10.18564/jasss.3212>

Baerenklau, K. A., & Provencher, B. (2005). Static modeling of dynamic recreation behavior: Implications for prediction and welfare estimation. *Journal of Environmental*

Economics and Management, 50(3), 617–636.

<https://doi.org/10.1016/j.jeem.2005.03.001>

Bailey, R. M., Carrella, E., Axtell, R., Burgess, M. G., Cabral, R. B., Drexler, M., ... Saul, S. (2018). A computational approach to managing coupled human–environmental systems: the POSEIDON model of ocean fisheries. *Sustainability Science*, 1–17.

<https://doi.org/10.1007/s11625-018-0579-9>

Bastardie, F., Nielsen, J. R., Andersen, B. S., & Eigaard, O. R. (2013). Integrating individual trip planning in energy efficiency - Building decision tree models for Danish fisheries. *Fisheries Research*, 143, 119–130.

<https://doi.org/10.1016/j.fishres.2013.01.018>

Bellman, M. A., & Heery, E. (2013). Discarding and fishing mortality trends in the U.S. west coast groundfish demersal trawl fishery. *Fisheries Research*, 147, 115–126.

<https://doi.org/10.1016/j.fishres.2013.04.007>

Bellman, M. A., Heppell, S., & Goldfinger, C. (2005). Evaluation of a US west coast groundfish habitat conservation regulation via analysis of spatial and temporal patterns of trawl fishing effort. *Canadian Journal of Fisheries and Aquatic Sciences*, 62(12), 2886–2900.

Boettiger, C., Mangel, M., & Munch, S. (2015). Avoiding tipping points in fisheries management through Gaussian process dynamic programming. *Proceedings of the*

Royal Society B: Biological Sciences, 282(1801), 20141631–20141631.

<https://doi.org/10.1098/rspb.2014.1631>

Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1), 5–32.

<https://doi.org/10.1023/A:1010933404324>

Carrella, E., Bailey, R. M., & Madsen, J. K. (2018). Repeated discrete choices in geographical agent based models with an application to fisheries. *Environmental Modelling & Software*. <https://doi.org/10.1016/J.ENVSOFT.2018.08.023>

Carrella, E., Bailey, R. M., & Madsen, J. K. (2018). Indirect inference through prediction. Retrieved from <http://arxiv.org/abs/1807.01579>

Cenek, M., & Franklin, M. (2017). An adaptable agent-based model for guiding multi-species Pacific salmon fisheries management within a SES framework. *Ecological Modelling*, 360, 132–149. <https://doi.org/10.1016/j.ecolmodel.2017.06.024>

Clark, C. W., & Mangel, M. (2000). *Dynamic state variable models in ecology : methods and applications* (1st ed.). New York: Oxford University Press.

Cooper, R., & Jarre, A. (2017). An Agent-based Model of the South African Offshore Hake Trawl Industry: Part I Model Description and Validation. *Ecological Economics*, 142, 268–281. <https://doi.org/10.1016/j.ecolecon.2017.06.026>

Csilléry, K., François, O., & Blum, M. G. B. (2012). Abc: An R package for approximate Bayesian computation (ABC). *Methods in Ecology and Evolution*, 3(3), 475–479.

<https://doi.org/10.1111/j.2041-210X.2011.00179.x>

Dorn, M. W. (2001). Fishing behavior of factory trawlers: A hierarchical model of information processing and decision-making. *ICES Journal of Marine Science*, 58(1), 238–252. <https://doi.org/10.1006/jmsc.2000.1006>

Dowling, N. A., Wilcox, C., Mangel, M., & Pascoe, S. (2012). Assessing opportunity and relocation costs of marine protected areas using a behavioural model of longline fleet dynamics. *Fish and Fisheries*, 13(2), 139–157.

<https://doi.org/10.1111/j.1467-2979.2011.00422.x>

Edmonds, B., & Moss, S. (2005). Multi-Agent and Multi-Agent-Based Simulation. *Multi-Agent and Multi-Agent-Based Simulation*, 3415, 130–144.

<https://doi.org/10.1007/b106991>

Elliston, L., & Cao, L. (2006). An agent-based bioeconomic model of a fishery with input controls. *Mathematical and Computer Modelling*, 44(5–6), 565–575.

<https://doi.org/10.1016/j.mcm.2006.01.010>

Evans, K. S., & Weninger, Q. (2014). Information Sharing and Cooperative Search in Fisheries. *Environmental and Resource Economics*, 58(3), 353–372.

<https://doi.org/10.1007/s10640-013-9701-8>

Gao, L., & Hailu, A. (2011). Evaluating the effects of area closure for recreational fishing in a coral reef ecosystem: The benefits of an integrated economic and biophysical modeling. *Ecological Economics*, 70(10), 1735–1745.

Gourieroux, C., Monfort, A., & Renault, E. (1993). Indirect Inference. *Journal of Applied Econometrics*, 8, 85–118. <https://doi.org/10.2307/2285076>

Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., ... DeAngelis, D. L. (2005). Pattern-Oriented Modeling of Agent Based Complex Systems: Lessons from Ecology. *Science*, 310(5750), 987–991.

Hartig, F., Calabrese, J. M., Reineking, B., Wiegand, T., & Huth, A. (2011). Statistical inference for stochastic simulation models - theory and application. *Ecology Letters*, 14(8), 816–827. <https://doi.org/10.1111/j.1461-0248.2011.01640.x>

Hastie, T., & Tibshirani, R. (2001). The elements of statistical learning (1st ed., Vol. 1). Berlin: Springer series in statistics Springer, Berlin. Retrieved from [http://crmtview.asc.ohio-state.edu/currofc/docs/115/Stat 760 Syllabus.pdf](http://crmtview.asc.ohio-state.edu/currofc/docs/115/Stat%20760%20Syllabus.pdf)

Haynie, A., & Layton, D. (2010). An expected profit model for monetizing fishing location choices. *Journal of Environmental Economics and Management*, 59(2), 165–176. <https://doi.org/10.1016/j.jeem.2009.11.001>

Hicks, A. C., & Wetzel, C. (2011). The status of Dover sole (*Microstomus pacificus*) along the US West Coast in 2011. Portland: Pacific Fishery Management Council.

Hicks, R. L., & Schnier, K. E. (2006). Dynamic random utility modeling: a Monte Carlo analysis. *American Journal of Agricultural Economics*, 88(4), 816–835.

<https://doi.org/10.2307/4123530>

Holland, D. S. (2008). Are Fishermen Rational? A Fishing Expedition. *Marine Resource Economics*, 23(3), 325–344. <https://doi.org/10.1086/mre.23.3.42629621>

Holland, D. S. (2016). Development of the Pacific Groundfish Trawl IFQ Market. *Marine Resource Economics*, 31(4), 453–464. <https://doi.org/10.1086/687829>

Holland, D. S., & Sutinen, J. G. (2000). Location Choice in New England Trawl Fisheries: Old Habits Die Hard. *Land Economics*, 76(1), 133.

<https://doi.org/10.2307/3147262>

Hunt, L. M. (2005). Recreational fishing site choice models: insights and future opportunities. *Human Dimensions of Wildlife: An International Journal*, 10(3), 153–172.

<https://doi.org/10.1080/10871200591003409>

Hunt, L. M., Boots, B. N., & Boxall, P. C. (2007). Predicting Fishing Participation and Site Choice While Accounting for Spatial Substitution, Trip Timing, and Trip Context. *North American Journal of Fisheries Management*, 27(3), 832–847.

<https://doi.org/10.1577/M06-079.1>

Hutniczak, B., & Münch, A. (2018). Fishermen's location choice under spatio-temporal update of expectations. *Journal of Choice Modelling*, 28, 124–136.

<https://doi.org/10.1016/j.jocm.2018.05.002>

Janssen, M. A. (2009). Understanding Artificial Anasazi. *Journal of Artificial Societies and Social Simulation*, 12(4), 13. <https://doi.org/13>

Kaplan, I. C., Holland, D. S., & Fulton, E. A. (2014). Finding the accelerator and brake in an individual quota fishery: Linking ecology, economics, and fleet dynamics of US West Coast trawl fisheries. *ICES Journal of Marine Science*, 71(2), 308–319.

<https://doi.org/10.1093/icesjms/fst114>

Keane, M. P., & Wolpin, K. I. (2007). Exploring the usefulness of a nonrandom holdout sample for model validation: Welfare effects on female behavior. *International Economic Review*, 48(4), 1351–1378. <https://doi.org/10.1111/j.1468-2354.2007.00465.x>

Klein, E. S., Barbier, M. R., & Watson, J. R. (2017). The dual impact of ecology and management on social incentives in marine common-pool resource systems. *Royal Society Open Science*, 4(8), 170740. <https://doi.org/10.1098/rsos.170740>

LeBaron, B. (2002). Short-memory traders and their impact on group learning in financial markets. *Proceedings of the National Academy of Sciences*, 99(Supplement 3), 7201–7206. <https://doi.org/10.1073/pnas.072079699>

Lian, C., Singh, R., & Weninger, Q. (2009). Fleet restructuring, rent generation, and the design of individual fishing quota programs: Empirical evidence from the Pacific Coast

groundfish fishery. *Marine Resource Economics*, 24, 329–359.

<https://doi.org/10.2307/42629661>

Lipsey, M. W. (1990). *Design Sensitivity: Statistical Power for Experimental Research*. SAGE.

Little, L. R., & McDonald, A. D. (2007). Simulations of agents in social networks harvesting a resource. *Ecological Modelling*, 204(3–4), 379–386.

<https://doi.org/10.1016/j.ecolmodel.2007.01.013>

Littman, M. L., Dean, T. L., & Kaelbling, L. P. (1995). On the complexity of solving Markov decision problems. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence* (pp. 394–402).

Marcoul, P., & Weninger, Q. (2008). Search and active learning with correlated information: Empirical evidence from mid-Atlantic clam fishermen. *Journal of Economic Dynamics and Control*, 32(6), 1921–1948. <https://doi.org/10.1016/j.jedc.2007.07.003>

Miller, J. H. (1998). Active Nonlinear Tests (ANTs) of Complex Simulation Models. *Management Science*, 44(6), 820–830.

Miller, S. J., & Deacon, R. T. (2017). Protecting Marine Ecosystems: Regulation Versus Market Incentives. *Marine Resource Economics*, 32(1), 83–107.

<https://doi.org/10.1086/689214>

Mistiaen, J. A., & Strand, I. E. (2000). Location choice of commercial fishermen with heterogeneous risk preferences. *American Journal of Agricultural Economics*, 82(5), 1184–1190. <https://doi.org/10.1111/0002-9092.00118>

Morton, M. (2005). Fishing for knowledge. *Australian Journal of Pharmacy*, 86(1020), 202. <https://doi.org/10.2307/3146679>

National Research Council (2006) *Dynamic Changes in Marine Ecosystems: Fishing, Food Webs, and Future Options*. National Academy Press, Washington, DC, pp. 154.

O'Hagan, A. (2006). Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering and System Safety*, 91(10–11), 1290–1300.
<https://doi.org/10.1016/j.ress.2005.11.025>

Ono, K., Haynie, A. C., Hollowed, A. B., Ianelli, J. N., McGilliard, C. R., & Punt, A. E. (2017). Management strategy analysis for multispecies fisheries including technical interactions and human behavior in modeling management decisions and fishing. *Canadian Journal of Fisheries and Aquatic Sciences*, cjfas-2017-0135.
<https://doi.org/10.1139/cjfas-2017-0135>

Parry, H. R., Topping, C. J., Kennedy, M. C., Boatman, N. D., & Murray, A. W. A. (2013). A Bayesian sensitivity analysis applied to an Agent-based model of bird population response to landscape change. *Environmental Modelling and Software*, 45, 104–115. <https://doi.org/10.1016/j.envsoft.2012.08.006>

Pudlo, P., Marin, J. M., Estoup, A., Cornuet, J. M., Gautier, M., & Robert, C. P. (2015). Reliable ABC model choice via random forests. *Bioinformatics*, 32(6), 859–866. <https://doi.org/10.1093/bioinformatics/btv684>

Reimer, M. N., Abbott, J. K., & Wilen, J. E. (2017). Fisheries Production: Management Institutions, Spatial Choice, and the Quest for Policy Invariance. *Marine Resource Economics*, 32(2), 143–168. <https://doi.org/10.1086/690678>

Richiardi, M., Leombruni, R., Saam, N., Sonnessa, M., Societies, A., & Simulation, S. (2006). A Common Protocol for Agent-Based Social Simulation. *Journal of Artificial Societies and Social Simulation*, 9(1). Retrieved from <http://jasss.soc.surrey.ac.uk/9/1/15.html>

Seri, R., & Secchi, D. (2017). How Many Times Should One Run a Computational Simulation? In B. Edmonds & R. Meyer (Eds.), *Simulating Social Complexity: A Handbook* (pp.229–251). https://doi.org/10.1007/978-3-319-66948-9_11

Schulze, J., Müller, B., Groeneveld, J., & Grimm, V. (2017). Agent-Based Modelling of Social-Ecological Systems: Achievements, Challenges, and a Way Forward. *Journal of Artificial Societies and Social Simulation*, 20(2), 8. <https://doi.org/10.18564/jasss.3423>

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., & De Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1), 148–175. <https://doi.org/10.1109/JPROC.2015.2494218>

Shelton, A. O., Thorson, J. T., Ward, E. J., & Feist, B. E. (2014). Spatial semiparametric models improve estimates of species abundance and distribution. *Canadian Journal of Fisheries and Aquatic Sciences*, 71(11), 1655–1666.

Simon, H. A. (1978). Rationality as process and as product of thought. *American Economic Association Review*, 68(2), 1–16. <https://doi.org/10.2307/1816653>

Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems* (pp. 1–12). Retrieved from <http://arxiv.org/abs/1206.2944>

Somers, K. A., Lee, Y. W., Jannot, J. ., & J., M. (2016). Catch tables by sector: Individual fishing quotas trawl and fixed gear, 2011-2015. Seattle, WA. Retrieved from http://www.nwfsc.noaa.gov/research/divisions/fram/observation/data_products/sector_products.cfm#obs

Steiner, E., Guldin, M., Warlick, A., & Pfeiffer, L. (2017). Economic Data Collection Program Catcher Vessel Report (2009-2015). Retrieved from https://www.nwfsc.noaa.gov/research/divisions/fram/documents/EDC_Catcher_Vessel_Report_June_2017.pdf

Stephens, A., & Taylor, I. G. (2013). Stock assessment and status of longspine thornyhead (*Sebastolobus altivelis*) off California, Oregon and Washington in 2013. Pacific Fisheries Management Council, Seattle, Washington, 123.

Stewart, I. J., Thorson, J. T., & Wetzel, C. (2011). Status of the US sablefish resource in 2011. Portland: Pacific Fishery Management Council.

Stow, C. A., Jolliff, J., McGillicuddy, D. J., Doney, S. C., Allen, J. I., Friedrichs, M. A. M., ... Wallhead, P. (2009). Skill assessment for coupled biological/physical models of marine systems. *Journal of Marine Systems*, 76(1–2), 4–15.

<https://doi.org/10.1016/j.jmarsys.2008.03.011>

Taylor, I. G., & Stephens, A. (2013). Stock Assessment of Shortspine Thornyhead in 2013. Portland: Pacific Fishery Management Council.

Taylor, I. G., & Wetzel, C. (2011). Status of the US yelloweye rockfish resource in 2011 (Update of 2009 assessment model). National Oceanic and Atmospheric Administration, National Marine Fisheries Service.

Toft, J. E., Punt, A. E., & Little, L. R. (2011). Modelling the economic and ecological impacts of the transition to individual transferable quotas in the multispecies US west coast groundfish trawl fleet. *ICES Journal of Marine Science*, 68(7), 1566–1579.

<https://doi.org/10.1093/icesjms/fsr095>

Uusitalo, L., Lehtikoinen, A., Helle, I., & Myrberg, K. (2015, January 1). An overview of methods to evaluate uncertainty of deterministic models in decision support.

Environmental Modelling and Software. Elsevier.

<https://doi.org/10.1016/j.envsoft.2014.09.017>

Van Putten, I. E., Kulmala, S., Thébaud, O., Dowling, N., Hamon, K. G., Hutton, T., & Pascoe, S. (2012). Theories and behavioural drivers underlying fleet dynamics models. *Fish and Fisheries*, 13(2), 216–235.

Windrum, P., Fagiolo, G., & Moneta, A. (2007). Empirical Validation of Agent Based Models : Alternatives and Prospects. *Journal of Artificial Societies and Social Simulation*, 10(2), 1–21. <https://doi.org/Article>

Zhao, L. (2010). A Model of Limit-Order Book Dynamics and a Consistent Estimation Procedure. Carnegie Mellon University.

Zou, H., & Hastie, T. (2005). Regression and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>