Rebuilding global fisheries under uncertainty

Milad Memarzadeh[a,1], Gregory L. Britten[b,c], Boris Worm[d], and Carl Boettiger[e,1]

[a]Department of Civil and Environmental Engineering, University of California, Berkeley, CA 94720; [b]Department of Earth System Science, University of California, Irvine, CA 92697; [c]Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139; [d]Department of Biology, Dalhousie University, Nova Scotia, B3H 4R2, Canada; and [e]Department of Environmental Science, Policy and Management, University of California, Berkeley, CA 94720

Edited by Stephan B. Munch, National Oceanic and Atmospheric Administration Fisheries, Santa Cruz, CA, and accepted by Editorial Board Member David W. Schindler June 24, 2019 (received for review February 14, 2019)

Current and future prospects for successfully rebuilding global fisheries remain debated due to uncertain stock status, variable management success, and disruptive environmental change. While scientists routinely account for some of this uncertainty in population models, the mechanisms by which this translates into decision-making and policy are problematic and can lead to unintentional overexploitation. Here, we explicitly track the role of measurement uncertainty and environmental variation in the decision-making process for setting catch quotas. Analyzing 109 well-sampled stocks from all oceans, we show that current practices may attain 55% recovery on average, while richer decision methods borrowed from robotics yield 85% recovery of global stocks by midcentury, higher economic returns, and greater robustness to environmental surprises. These results challenge the consensus that global fisheries can be rebuilt by existing approaches alone, while also underscoring that rebuilding stocks may still be achieved by improved decision-making tools that optimally manage this uncertainty.

fisheries | decision theory | adaptive management

> Managing fisheries is hard: it's like managing a forest, in which the trees are invisible and keep moving around.
>
> John Shepherd, circa 1978

Previous controversy over the future of global fisheries (1–3) is gradually giving way to a growing consensus on solutions: in overcoming historical overexploitation in global fisheries, reforms based on existing methods will be sufficient to secure the recovery of most stocks (4–6). This optimism may be premature. Fisheries management is complicated by profound uncertainty over the number of fish that are in the sea—even well-designed scientific surveys and sophisticated models leave us with large measurement errors regarding the state of a particular stock. This is not the only limiting assumption made in models forecasting the recovery of fish stocks, but it is uniquely problematic for current decision methods that translate model output into management advice. Here, we revisit the role of measurement uncertainty in fisheries management by analyzing populations using decision methods that allow us to fully propagate uncertainty due to imperfect estimates of population size through future population trajectories and harvest scenarios.

Most decision methods currently used in fisheries are attempting to maximize some objective such as maximum sustainable yield (MSY) or maximum economic yield (MEY). While the underlying fish-stock assessment models typically account for measurement errors (7), the optimization procedure that is used in decision-making often fails to account for that uncertainty. For example, MSY is based on a static optimization that seeks to determine a constant mortality, implicitly ignoring both stochasticity in stock dynamics and uncertainty in measurement. This approach is rooted in foundations of fisheries management science (8, 9) and remains the operating principle in many current international fishing agreements, including the United Nations Law of the Sea (10).

Dynamic optimization techniques have emerged more recently, showing that in stochastic environments, varying mortality targets can produce higher yields than the constant mortality rule of MSY (see formal proof in ref. 11). Because dynamic approaches must be defined in terms of discounted economic value, they have been favored more by resource economists and have been used in determining MEY (5), or in the context of rights-based fisheries management (RBFM) (4). The underlying solution method uses Markov decision processes (MDPs) (12). The Markov property dictates that we can know the probability of any future state of the system if we only know the current state (given the present, the future is independent of the past). It has long been recognized that this property does not extend to imperfect observations of the state variable, whose probabilities will depend on all previous observations (13–15). Such decision problems belong to a class of problems known as partially observable MDPs (POMDPs), which are not amenable to the same algorithms and have so far not been solved in a fisheries context (13, 16). Meanwhile, recognition that such optimization-based approaches ignore the reality of imperfect measurements, managers frequently seek heuristic adjustments. Management strategy evaluation (MSE) compares the performance of a given set of alternate policies through forward simulations (17). While MSE provides a mechanism to compare among policies derived from approaches such as MSY, MEY, or heuristic adjustments of those rules such as pretty good yield (PGY) (18), it can only evaluate proposed strategies, not generate new ones.

Using an MSE approach (17), we evaluate the performance of existing decision methods, finding that under sufficiently large

Significance

Many fisheries that have been historically overexploited are now considered to be rebuilding, with a hope that current best practices could ensure the recovery of most overfished species by mid century. Our analysis suggests this optimism may be premature, as current projections typically assume managers can have perfect measurements of current stock sizes. We demonstrate how such an assumption can undermine rebuilding efforts under current best practices and even drive unintentional stock declines. By borrowing novel decision methods from the field of robotics, we also show how stock rebuilding can be achieved in the face of measurement and environmental uncertainty, while also achieving higher economic returns than expected under current approaches.

Author contributions: M.M. and C.B. designed research; M.M., G.L.B., B.W., and C.B. performed research; M.M. and C.B. contributed new reagents/analytic tools; M.M. and C.B. analyzed data; and M.M., G.L.B., B.W., and C.B. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. S.B.M. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: memarzadeh.milad@gmail.com or cboettig@berkeley.edu.

Published online July 22, 2019.

measurement errors (10 to 20%), these methods may impede recovery or cause net declines even in carefully managed stocks where professional stock assessment estimates are available. We then take advantage of relatively recent advances in the fields of robotics and autonomous vehicle navigation (19), which have spurred the development of state-of-the-art algorithms making large POMDP problems tractable (20), to solve for the optimal strategy under observation uncertainty in the fisheries context. We demonstrate that POMDP-derived policies could still provide for the recovery of these stocks by midcentury.

Materials and Methods

We examine 4 approaches to future fishery management: 1) business-as-usual management (BAU), in which current fishing effort is used for projections; 2) maximizing harvest (MSY); 3) maximize long-term economic value, assuming the stock size is known (MDP); and 4) maximizing economic value when stock size is uncertain (POMDP). For each fishery, we estimate catch, profit, and biomass under each approach from today to 2050. To simplify the comparison between economic maximization approaches (MDP and POMDP) and catch maximization approaches (MSY), we use a trivial economic function in which profits are directly proportional to catch. Other societal objectives such as employment, equity, or biodiversity conservation are clearly important but are not explicitly modeled here. While it is possible to include both multiple objectives and price dynamics in the profit function (21), this can exaggerate the performance of economic-based optimization.

We consider the population dynamics of a generic fishery as follows:

$$b_t = \zeta_t g\left(b_{t-1}, q_{t-1}\right)$$
$$z_t = \epsilon_t b_t,$$

where b_t and q_t are the population biomass and catch quota at timestep (i.e., year) t, g describes the expected population dynamics, ζ_t captures the inherent stochasticity in population growth, and z_t is the estimated stock size subject to measurement error ϵ_t. The decision methods based on static maximization of the long-term catch (i.e., MSY) ignore both ζ_t and ϵ_t, while existing dynamic optimizations (i.e., MDP) (4, 5) ignore ϵ_t. Here, we implement POMDPs, which simultaneously model and incorporate both sources of uncertainty.

It is important to remember that the decision problem takes the population dynamics model as a given input and seeks to determine the policy which maximizes some specified objective—in our case, economic value from harvesting. For simplicity, we assume the commonly used Gordon–Schaefer surplus production model for g (22). We allow both random terms to follow a Gaussian distribution with ζ_t and ϵ_t mean unity and standard deviations of σ_g and σ_m, respectively (see ref. 23). While management of individual fisheries frequently employs more complex population models reflecting aspects such as age structure in the population, data limitations lead global analyses to favor simpler approximations such as Gordon–Schaefer (4). As the methods considered here (MSY, MDP, POMDP) are applicable to any population dynamics model, controlling for the model across all methods allows for a direct comparison both between the methods and against previous global analyses (4).

Alternate formulations of this decision problem are also possible, such as more or less frequent observations, or policies formulated in terms of fishing effort or landing fees rather than quotas (24). All of the decision methods compared here are agnostic to these details, just as they are agnostic to the choice of population dynamics. That is, one could also use MSY, MDP, or POMDP to decide on a fishing effort rather than a quota. Nevertheless, our focus on quota-based policies is not arbitrary. While it is possible to regulate some aspect of fishing effort directly (length of season, size of vessel), effort or mortality-based targets are frequently implemented in terms of catch quotas (e.g., by setting a total allowable catch [TAC]). Using quota-based policies also allows direct comparison with (4), which ignores measurement uncertainty.

Fishing quotas are determined as follows. In the case of BAU, the management quota is defined as the last observed fishing mortality $H_t = F_0 B_t$, where F_0 is the fishing mortality last observed in that stock and B_t is the estimated biomass. The MSY quota is defined similarly but with mortality fixed to MSY rather than historical observation, $H_t = F_{MSY} B_t$. Under MDP and POMDP, quotas are determined directly by the corresponding algorithms. These definitions are consistent with how the same policies are defined in ref. 4. To simplify comparisons across stocks of widely varying total biomass, we also measure recovery in terms of individual stock biomass relative to 80% of the stock's estimated biomass predicted to achieve MSY, B_{MSY}.

We explicitly model the dynamics for each fishery under each harvest policy. To do so, we estimated the intrinsic growth rate and carrying capacity of each fishery in the dataset. We consider forecasts for 109 commercially harvested marine fisheries for which sufficient data are available (over 30 contiguous data points; Dataset S1) in the R. A. Myers (RAM) Legacy Stock Assessment Database v3.0 (25) for reliable model estimates (*SI Appendix*, *Materials and Methods*). For each stock, we perform a Bayesian estimation of parameters of the population dynamics model for the stochastic Gordon–Schaefer recruitment model (*SI Appendix*, *Materials and Methods* and Dataset S2). We then consider 500 replicate simulations of the estimated stock dynamics.

Estimating both measurement uncertainty and intrinsic stochasticity (environmental and demographic noise) (23) from the RAM Legacy Stock Assessment data can be subtle. The assessment data are themselves the result of model output, sometimes averaged across multiple models. Consequently, this means that our estimates using the RAM data may underestimate intrinsic stochasticity, especially in those regions where RAM data appear to be only smooth model hindcasts. Since higher levels of intrinsic stochasticity will tend to improve the performance of the POMDP relative to existing approaches, underestimating the noise would only favor the simpler existing methods. Estimating just how uncertain stock size estimates are can be more difficult than estimating the stock size itself. Deviations between true biomass and that estimated by a stock assessment arise in many ways, and accepted ranges of uncertainty vary greatly (7). We address the issue by repeating our analyses over a range of possible measurement errors (0%, 5%, 10%, 20%; for details, see *SI Appendix*, *Materials and Methods*). While the upper bound of uncertainty is often put even higher than 20% in some stocks, this range is sufficient to see the consequences of both small and larger measurement error.

Results and Discussion

We find that current decision methods fail to rebuild many stocks under moderate uncertainty. Fig. 1 shows the percentage of global stocks in our analysis that remain above or are rebuilt to the threshold of 80% B_{MSY} over time under each decision method. In the absence of measurement error, we see a broad long-term recovery of stocks toward their target biomass, B_{MSY} by 2050, reversing declines seen in the previous half-century (Fig. 1*A*), consistent with previous findings (4). In this scenario, MDP and POMDP solutions are mathematically identical and achieve a much faster rate of recovery than MSY and BAU. Due to stochastic stock recruitment (e.g., environmental noise), some highly overexploited stocks are lost before they have a chance to recover. This fraction is higher under the MSY and BAU decision methods (as shown in Fig. 1*A*, the recovery rate of POMDP and MDP is significantly higher with perfect measurement of population size).

The introduction of increasingly severe measurement error (5%, 10%, 20%; Fig. 1 *B*–*D*, respectively), impedes stock recovery under MSY, BAU, and MDP decision methods, while the POMDP solution proves robust to these errors. The MDP decision method (used to solve the RBFM scenario of ref. 4) is particularly vulnerable to large measurement errors, achieving a lower recovery rate than the other methods at 10% and driving outright decline at 20%. Although the majority of stocks recover under MSY at all levels of uncertainty, this fraction is considerably smaller than projected under a POMDP method.

Patterns of recovery vary by individual region and species. Fig. 2 compares projected recovery rates to the year 2050 assuming a measurement error of 10% by region and species group. In the absence of this error, nearly all overexploited species recover to above 80% B_{MSY} (*SI Appendix*, Fig. S1). Under a 10% measurement error, Fig. 2 shows that MDP and MSY decision methods see many stocks failing to recover, with the impact varying considerably by region or species, owing to differences in current status and estimated population dynamics between different stocks (Fig. 2*B*). The MDP decision method performs sometimes better and sometimes worse than MSY, while the POMDP decision method permits strong rebuilding of stocks across all regions and species groups.

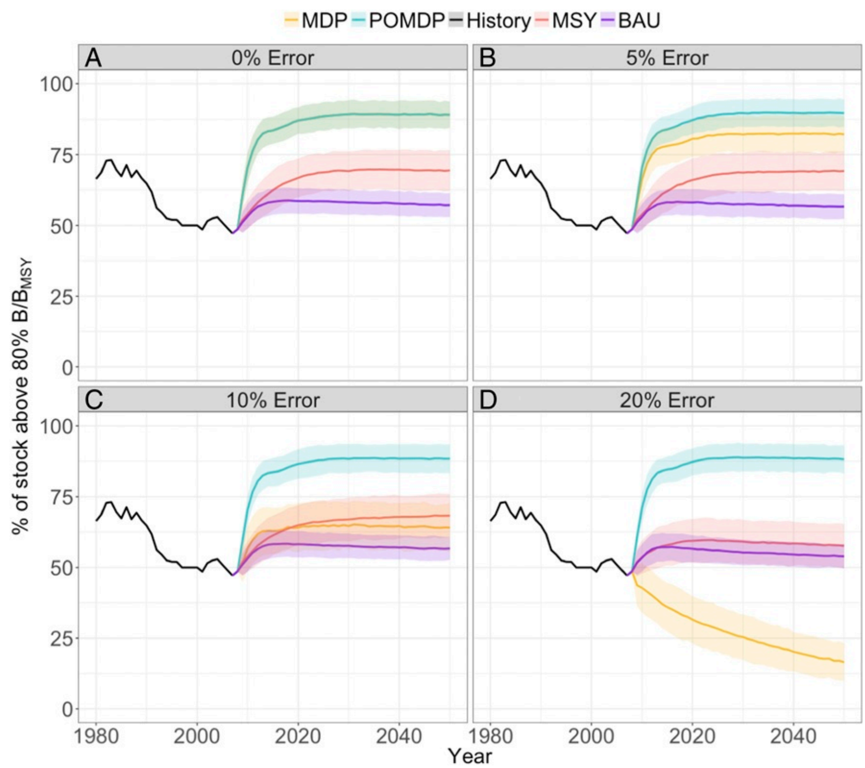

Fig. 1. Projected rebuilding of global fish stocks. Historical trajectories and projections of stocks being above 80% B_{MSY} under different management regimes. Black indicates historical observations; blue, POMDP; yellow, MDP; red, MSY; and purple, BAU. Different graphs correspond to the intensity of the present measurement error in estimating the population biomass, i.e., 0% (*A*), 5% (*B*), 10% (*C*), and 20% (*D*). Lines show average trends over all stocks listed in Dataset S1, and shades show ±SD. It should be noted that in *A*, MDP (yellow) and POMDP (blue) solutions are mathematically identical and result in the green curve.

A closer look at the replicate simulations for several representative stocks can provide a clearer picture of how current decision methods may fall short of rebuilding global fisheries. Under measurement uncertainty, some stocks merely recover more slowly under MSY or achieve lower long-term biomass under MDP than under the optimal management provided by POMDP-based decision-making (e.g., New Zealand Black cardinalfish, *Epigonus telescopus* [Fig. 3*A*], a slow growing species). In other stocks, particularly those with faster growth rates, the introduction of imperfect measurements in a stochastic environment make projected rebuilding under MSY highly variable (e.g., in South Pacific Horse mackerel, *Trachurus trachurus*; Fig. 3*B*). In some cases, this variation is sufficient to drive long-term declines in the expected stock size under MSY-based management (Atlantic cod, *Gadus morhua*, Scotian shelf; Fig. 3*C*). While MDP-based management is less vulnerable to heavy overfishing under imperfect measurements, these errors can still be sufficient to prevent recovery to 80% of B_{MSY} and can, in some cases, still result in long-term average declines in predicted biomass (e.g., Pacific herring, *Culpia pallassi*, Straight of Georgia [Fig. 3*D*]; for the corresponding catch values, refer to *SI Appendix*, Fig. S3).

Together, these results show that individual life-history differences can provide further insight into which species are most impacted by the uncertainty. Slow-growing species like the Black cardinalfish rebuild particularly slowly under MSY-based management, but MDP or MSY management may be most problematic for faster-growing species with strong coupling to environmental fluctuations, like the Pacific herring. However, the relative performance of the decision methods is consistent across these differences in life history and environment. Differences in the magnitude of measurement errors have a larger impact, with MSY generally doing worse with measurement error $\leq 10\%$ and MDP doing worse with larger error.

SI Appendix, Fig. S2 demonstrates that the consistently higher rate of fish stock recovery across species and regions under a POMDP-based decision method does not come at the cost of reduced economic returns. Rather, ignoring uncertainty also reduces the long-term economic value of stocks, with MSY achieving 65 to 80% of the value realized under the POMDP decision method, and MDP achieving 20 to 80% (*SI Appendix*, Fig. S2). Overexploited stocks quickly become less productive as biomass falls below B_{MSY}, making overharvesting as well as underharvesting more economically costly. The dynamic MDP decision method is particularly sensitive to measurement error. The reason lies in the Markovian assumption. Under imperfect measurement of population size, management of future of stocks needs to consider all historical observations (14). MDP ignores this complexity, which results in a poor performance in the presence of the measurement error.

The greater performance of POMDP in terms of stock recovery rates is perhaps more surprising than the performance in terms of economic value. For simplicity and consistency with previous analyses (4), we have used each of the decision methods here to try and maximize a simple measure of economic value based only on catch. While economic optimization methods such as MDP and POMDP can be adapted to more general notions of economic value (25), to do so would bias the comparison against MSY. Because POMDP approach solves for the optimal strategy under uncertainty, it by definition sets the bar for economic performance in *SI Appendix*, Fig. S2. However, since the economic objective used here merely maximizes long-term catch, it is in no way assured that doing so will also rebuild by midcentury most

Fig. 2. Regional and species-specific patterns of stock rebuilding. Bar charts show the fraction of stocks that have recovered to above 80% B_{MSY} by 2050, which are currently estimated as below that target. *A* shows projected outcomes for stocks grouped by geographic region (refer to Dataset S3), *B* shows outcomes for stocks grouped by species (refer to Dataset S1). Projections show averages over 500 replicate simulations at 10% measurement error.

of the stocks we have examined. This ultimately depends on the discount rate applied to future economic returns. When the value of future harvests is very highly discounted, the optimal strategy may harvest the stock to extinction. Our analysis has assumed a modest annual discount rate of 1%, which rewards sustaining yields over the long term. Alternate discount rates are compared in *SI Appendix*, Fig. S4.

Given the difficulty in accurately estimating uncertainty even in the RAM data, our analysis has considered a plausible range of measurement error between 0 and 20% and found that the relative comparisons between different decision methods are consistent across this range. Without perfect information, we find MSY would achieve only 10 to 30% recovery of overexploited stocks, while MDP would achieve between 0 to 60% recovery (Fig. 1). We implement a POMDP-based decision method that efficiently incorporates uncertainty in observed population size, allowing over 80% of overexploited stocks to recover by midcentury even under substantial uncertainty (Fig. 1*D*), while also projecting significantly higher economic returns than expected under current decision methods after accounting for this uncertainty (*SI Appendix*, Fig. S2).

Why do existing decision methods perform so poorly under measurement uncertainty, given that this error is equally likely to underestimate as to overestimate the current stock size? The intuition for this result was anticipated by Clark and Kirkwood (13), who observed that under traditional decision methods, there is never any projected risk of local stock collapse and extinction regardless of the uncertainty in future growth rates. The frequency of local stock extinctions across simulations varies with measurement error and management strategy. At 10% measurement error, our simulations find an average of 6% of all stocks are expected to experience collapse within 50 y under the MDP-based decision method and 17% collapse under MSY (compared with 4% under POMDP). Local extinction arises from the interaction of stochastic growth and measurement error. Under our model, stochastic extinction is possible but only likely to be observed from very low stock sizes. MSY-based management is more susceptible to extinction at moderate measurement error, where a sequence of overestimations leads to overharvesting even from low stock sizes. Because the MDP strategy does not harvest at all at very low stock sizes, it is less susceptible to such moderate measurement error. MDP is more sensitive to large measurement errors, since the "bang-bang" nature of its constant-escapement solution harvests in excess of MSY when stock sizes exceed B_{MSY}. Because BAU is implemented like MSY as a constant effort rather than a constant escapement, it too is less sensitive to measurement error, although vulnerable to intrinsic stochasticity.

Our results imply that current approaches to decision-making (including MSY- and MDP-based methods) are not sufficient to rebuild global fisheries. In their place, newer and more rigorous but computationally intensive approaches such as POMDP may prove vital to future conservation planning and resource management. There is some irony in the fact that fisheries applications spurred early development of such methods (26), but sophisticated algorithms such as those we have used here have been largely developed in other fields (20, 27, 28). POMDP decision methods have proven promising in other conservation applications as well (29, 30). It is time to return them to fisheries.

Our conclusions are not without caveats. First, our analysis has focused only on stocks for which scientific assessments of stock biomass were readily available (*SI Appendix*). Although this represents a small fraction of commercially exploited stocks worldwide, these are some of the better-managed and least-uncertain stocks (4), which suggests that current decision methods would fare even more poorly with other stocks, making our conclusions conservative. Second, our results have followed other global analyses in assuming a simple model of population dynamics.

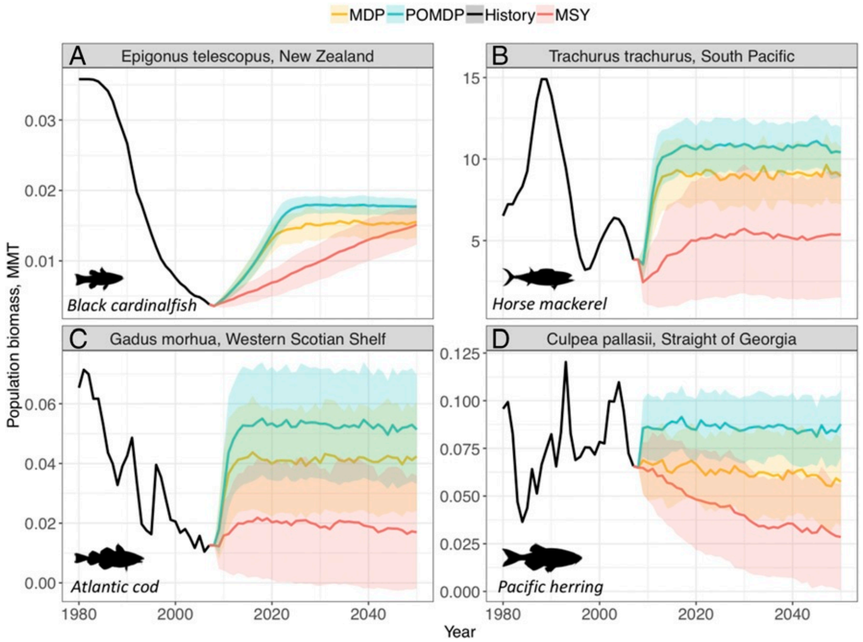

Fig. 3. Example projections of individual fish population. This figure shows 4 example population biomass of Black cardinalfish (*E. telescopus*) (*A*), Horse mackerel (*T. trachurus*) (*B*), Atlantic cod (*G. morhua*) (*C*), and Pacific herring (*Culpea pallasii*) (*D*) and their projections based on different decision methods. Projections show averages over 500 replicate simulations at 10% measurement error. Lines show average trends over replicates, and shades show ±SD. For the corresponding catch values, refer to *SI Appendix*, Fig. S3.

Extensions to stage-structured models may be particularly fruitful where fisheries management weighs short- versus long-term benefits of preserving population structure while accounting for uncertainty in 100, or more, model parameters. All of the decision methods considered here (MSY, MDP, POMDP) can, in principle, be applied to more complex stage-structured models, although at greater computational cost. Further work is also needed to address nonstationary models, in which parameters may change over time due to external forces, such as climate change, as well as autocorrelation in environmental fluctuations. Finally, decision methods can only be as good as the objectives they are given. More work is needed to properly incorporate additional objectives such as employment, equity, or biodiversity conservation into these approaches. Even so, a decision method provides only recommendations on which to base a policy, and significant challenges remain to implementing, enforcing, and adjusting such policies. Many small-scale fisheries are not managed at all, and underreporting catch, bycatch, and discards are an important problem in larger fleets.

Ecologists have for some time recognized the growing divide between the sophistication of population models and the limitations of decision theory through which they are applied in resource management, and some have called for new approaches to bridge this gap (31, 32). Our approach provides an important step in this direction. Of course, any decision method is only as good as the human institutions that make and implement those policies, but so too should those policies reflect our best methods. As the rebuilding of depleted fisheries is becoming a unifying policy goal, outdated concepts must give way to more modern approaches that take full count of substantial and growing uncertainty we face in the ocean.

ACKNOWLEDGMENTS. C.B. acknowledges computational resources from Extreme Science and Engineering Discovery Environment Jetstream (NSF Division of Environmental Biology Grant 160003) and Chameleon Testbed, supported by NSF Grants 1419141 and 1419152, as well as the support by the US Department of Agriculture Hatch project (Grant CA-B-INS-0162-H). G.L.B. was supported by the Simons Collaboration on Computational Biogeochemical Modeling of Marine Ecosystems (Grant 549931).

1. B. Worm *et al.*, Impacts of biodiversity loss on ocean ecosystem services. *Science* **314**, 787–790 (2006).
2. B. Worm *et al.*, Rebuilding global fisheries. *Science* **325**, 578–585 (2009).
3. T. A. Branch, O. P. Jensen, D. Ricard, Y. Ye, R. Hilborn, Contrasting global trends in marine fishery status obtained from catches and from stock assessments. *Conserv. Biol. J. Soc. Conserv. Biol.* **25**, 777–786 (2011).
4. C. Costello *et al.*, Global fishery prospects under contrasting management regimes. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 5125–5129 (2016).
5. M. G. Burgess *et al.*, Protecting marine mammals, turtles, and birds by rebuilding global fisheries. *Science* **359**, 1255–1258 (2018).
6. G. L. Britten, M. Dowd, L. Kanary, B. Worm, Extended fisheries recovery timelines in a changing environment. *Nat. Commun.* **8**, 15325 (2017).
7. C. M. Dichmont *et al.*, A review of stock assessment packages in the United States. *Fish Res.* **183**(Suppl C), 447–460 (2016).
8. M. B. Schaefer, Some aspects of the dynamics of populations important to the management of the commercial marine fisheries. *Bull. Inter Am. Trop. Tuna Comm.* **1**, 27–56 (1954).
9. H. Scott Gordon, The economic theory of a common-property resource: The fishery. *J. Polit. Econom.* **62**, 124–142 (1954).
10. United Nations, "United Nations Convention on the Law of the Sea" (Tech. Rep. 1833 UNTS 3, United Nations, Geneva, Switzerland, 1982).
11. W. J. Reed, Optimal escapement levels in stochastic and deterministic harvesting models. *J. Environ. Econ. Manage.* **6**, 350–363 (1979).
12. C. W. Clark, Profit maximization and the extinction of animal species. *J. Polit. Econom.* **81**, 950–961 (1973).
13. C. W. Clark, G. P. Kirkwood, On uncertain renewable resource stocks: Optimal harvest policies and the value of stock surveys. *J. Environ. Econ. Manage.* **13**, 235–244 (1986).
14. M. Memarzadeh, C. Boettiger, Resolving the measurement uncertainty paradox in ecological management. *Am. Nat.* **193**, 645–660 (2019).
15. M. Memarzadeh, C. Boettiger, Adaptive management of ecological systems under partial observability. *Biol. Conserv.* **224**, 9–15 (2018).
16. E. J. Sondik, The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs. *Oper. Res.* **26**, 282–304 (1978).
17. A. E. Punt, D. S. Butterworth, C. L. de Moor, J. A. A. De Oliveira, M. Haddon, Management strategy evaluation: Best practices. *Fish Fish.* **17**, 303–334 (2016).
18. R. Hilborn, Pretty good yield and exploited fishes. *Mar. Policy* **34**, 193–196 (2010).
19. H. Kurniawati, D. Hsu, W. Sun Lee, "SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces" in *Proceedings of Robotics: Science and Systems IV*, O. Brock, J. Trinkle, F. Ramos, Eds. (MIT Press, Cambridge, Massachusetts, 2009).
20. N. Roy, G. Gordon, S. Thrun, Finding approximate POMDP solutions through belief compression. *J. Artif. Intell. Res.* **23**, 1–40 (2005).

21. B. S. Halpern *et al.*, Achieving the triple bottom line in the face of inherent trade-offs among social equity, economic return, and conservation. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 6229–6234 (2013).
22. M. B. Schaefer, Some considerations of population dynamics and economics in relation to the management of the commercial marine fisheries. *J. Fish. Res. Board Can.* **14**, 669–681 (1957).
23. C. Boettiger, From noise to knowledge: How randomness generates novel phenomena and reveals information. *Ecol. Lett.* **21**, 1255–1267 (2018).
24. M. L. Weitzman, Landing fees vs harvest quotas with uncertain fish stocks. *J. Environ. Econ. Manage.* **43**, 325–338 (2002).
25. D. Ricard, C. Minto, O. P. Jensen, J. K. Baum, Examining the knowledge base and status of commercially exploited marine species with the RAM legacy stock assessment database. *Fish Fish.* **13**, 380–398 (2011).
26. R. E. Bellman, *Dynamic Programming* (Princeton Univ. Press, Princeton, 1957), pp. 317–339.
27. M. Memarzadeh, M. Pozzi, J. Z. Kolter, Optimal planning and learning in uncertain environments for the management of wind farms. *J. Comput. Civil Eng.* **29**, 04014076 (2014).
28. M. Pozzi, M. Memarzadeh, K. Klima, Hidden-model processes for adaptive management under uncertain climate change. *J. Infrastruct. Syst.* **23**, 04017022 (2017).
29. I. Chadès *et al.*, When to stop managing or surveying cryptic threatened species. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 13936–13940 (2008).
30. P. Fackler, K. Pacifici, Addressing structural and observational uncertainty in resource management. *Environ. Manage.* **133**, 27–36 (2014).
31. J. Fischer *et al.*, Integrating resilience thinking and optimisation for conservation. *Trends Ecol. Evol.* **24**, 549–554 (2009).
32. S. Polasky, S. R. Carpenter, C. Folke, B. Keeler, Decision-making under great uncertainty: Environmental management in an era of global change. *Trends Ecol. Evol.* **26**, 398–404 (2011).