



Identifying fish populations prone to abrupt shifts via dynamical footprint analysis

Alejandro V. Cano^{a,1} , Olaf P. Jensen^b , and Vasilis Dakos^a

Edited by Alan Hastings, University of California, Davis, CA; received March 11, 2025; accepted July 11, 2025

Fish population biomass fluctuates through time in ways that may be either gradual or abrupt. While abrupt shifts in fish population productivity have been shown to be common, they are rarely integrated into stock assessment or fishery management, in part because of the difficulty of predicting when abrupt shifts may occur and which stocks are prone to such shifts. In this study, we address the latter challenge by designing a mechanism-agnostic context-specific approach that is based on exploiting the dynamical properties of fish population fluctuations for detecting potential abrupt shifts. We use time series of fish population biomass from three global datasets, first, to classify their shapes into abrupt and nonabrupt (linear, quadratic, or no change) classes, and, second, to predict classified shapes based only on their dynamical footprint (a set of metrics such as variance, autocorrelation, etc., of the time series). We find that populations prone to abrupt shifts can be detected with moderate accuracy in the three datasets in spite of data limitations. In total, we identified 50 populations at risk of future abrupt shifts across 11 different Large Marine Ecosystem regions. Our context-specific approach offers critical insights into population stability and enables the identification of stocks whose dynamical properties suggest that they would benefit from more precautionary management.

stability | population dynamics | regime shifts | machine learning

Despite recent improvements in the sustainable management of marine fisheries (1), many remain overfished (2, 3). A variety of management measures show promise for reducing overfishing (4). Currently, however, fish populations are increasingly vulnerable to multiple anthropogenic pressures, creating urgent challenges for both biodiversity conservation and the livelihoods of communities that depend on marine resources (5–7). For example, climate change, habitat degradation, and other direct and indirect stressors (8, 9) combined with fishing are responsible for the collapse of some of the most historically plentiful fisheries (10). In many cases, such collapses occur in sudden, abrupt ways (11) leading to irreversible shifts that are hard-to-predict (12). With such increasing pressure on fish populations worldwide, context-specific assessments of the risk of abrupt shifts in fish populations can help support their long-term sustainability.

Although abrupt shifts (cf regime shifts) in marine fish populations are well documented for a variety of cases, there is less work on their understanding at large scales (13, 14) and even less on evaluating their risk of occurrence. Most work is based on the detailed analysis of the drivers of regime shifts in iconic case studies, like the North Atlantic cod (Keith), or at community scale shifts, like the Baltic (15) and North Seas (16). Fish population modeling (i.e., stock assessment) has traditionally utilized models with smooth or continuous population changes to determine sustainable harvest levels and identify populations experiencing overfishing (17). Only relatively recently, have dynamical properties based on the fluctuations of statistical moments of time series been used to detect the abrupt collapse or recovery of marine populations (18–21). Such dynamical properties are considered generic proxies of a change in the dynamical stability of a population and their theoretically expected patterns over time can be used to infer upcoming abrupt shifts in time to take appropriate management action (22).

In this study, rather than assuming a priori expected dynamical stability patterns in populations experiencing abrupt shifts, we adopt an empirical context-specific approach. Specifically, we 1) categorize each fish population time series by identifying its characteristic shape (abrupt, quadratic, linear, and no change), 2) compute the dynamical footprint of each population (i.e., an ensemble of stability-related metrics), and 3) use a boosted random tree model (BRT) to predict the shape classes using their dynamical footprint (Fig. 1). This approach allows us to evaluate whether the dynamical patterns of fish populations differ between populations that have or have not experienced an abrupt shift and ultimately to help identify potential future abrupt shifts in their dynamics.

Significance

The stability of fish populations is crucial for food security and biodiversity conservation. Fish populations can experience abrupt and unexpected changes over time, and understanding when such changes could occur is invaluable for sustainable management. Here, by using machine learning and analyzing the shapes and patterns underlying fish population dynamics from three global datasets, we identify populations prone to abrupt shifts that may require management intervention. Our approach could be adapted to other ecosystems fostering new developments in the estimation of ecological resilience and targeted conservation measures.

Author affiliations: ^aInstitut des Sciences de l'Évolution de Montpellier, Univ. de Montpellier, CNRS, Institut de recherche pour le développement, Montpellier 34090, France; and ^bCenter for Limnology, University of Wisconsin-Madison, Madison, WI 53706

Author contributions: A.V.C., O.P.J., and V.D. designed research; A.V.C. performed research; A.V.C. analyzed data; and A.V.C., O.P.J., and V.D. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2025 the Author(s). Published by PNAS. This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: alejandro.vilorio-cano@umontpellier.fr.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2505461122/-DCSupplemental>.

Published August 18, 2025.




Fig. 1. Workflow of dynamical footprint analysis for abrupt shift detection. We employ a 5 step workflow on three different curated datasets of fish population time series: 1) productivity time series from the RAM Legacy dataset, and CPUE time series for 2) the ICES Atlantic Cod and 3) FishGlob datasets (*Materials and Methods*). (step A) Following ref. 14, we classify time series according to their shapes into four classes: abrupt, quadratic, linear, and no change. (step B) We detrend the time series using the best shape model found in (step A), and if the time series is classified as “abrupt” we trim the time series to preserve the interval before the detected shift for the next steps. (step C) The dynamical footprint of each time series is composed of 7 metrics (DEV, ar1, SD, dispersion, disparity, CV, skewness) that are calculated in two ways: using the full time series (metric) and using a rolling window (metric_trend) (*Materials and Methods*). (step D) We train a gradient boosted random tree model (BRT) to predict the shape class based on the dynamical footprint estimated in (step C). (step E) Finally, we compare the actual shape classes given by the classification in (step A) to the shape classes predicted by the BRT model in (step D). The resulting matrix presents aggregated results for two classes: “abrupt” and “nonabrupt.” A “potential abrupt” population is also identified, representing cases initially classified as “nonabrupt” based on their shape. However, these cases exhibit a dynamical footprint that the BRT model recognizes as sharing characteristics with “abrupt” shifts.

We apply this approach to fish population time series from three different databases: The RAM Legacy (23), ICES Atlantic Cod (24), and FishGlob (25) databases. RAM Legacy provides comprehensive global stock assessment model output (estimated biomass) and associated catch time series for predominately commercially exploited marine species. The ICES Atlantic Cod database is a collection of trawl survey data (catch per unit of effort) for surveys targeting Atlantic cod in European waters. FishGlob is also a collection of trawl survey data, but offers a broader perspective by integrating multiple regional surveys. Together, these databases can be used to empirically characterize the dynamics of fish populations and could enhance our understanding of fish population stability and their potential risk to abrupt shifts, albeit from different perspectives. Despite the inherent heterogeneity within each dataset, we hypothesize that each dataset still retains common unique idiosyncrasies (e.g., sampling method, dataset curation, unit, species, management, etc), which should be reflected in distinct dynamical footprints between datasets. This is what we call an empirical context-specific approach as we analyze each dataset individually to preserve their distinctive features.

Results

Shape Classification of Fish Population Time Series. The three fish populations time series datasets contained 9,006 time series in total, 397 productivity (RAM Legacy), 199 catch per unit of effort (CPUE) (ICES cod), and 8410 CPUE (FishGlob). After discarding time series with missing values, as well as the ones with length below 20 y (see *Materials and Methods* for further details), we classified the remaining 629 time series into four classes based on their shape: (14) no change, linear, quadratic, and abrupt

(*Materials and Methods*). All four shape classes were found in the three datasets with different relative proportions (Fig. 1A and *SI Appendix, Fig. S1*). While “quadratic” time series were predominant in both the RAM Legacy and the FishGlob datasets, for the ICES cod dataset most time series were classified as “no change.” Among populations classified as “abrupt,” the ICES cod dataset exhibited the largest proportion with a decreasing trend (50%), while RAM Legacy and FishGlob showed 24% and 28%, respectively. We found that classification quality scores were higher for “abrupt” and “quadratic” time series (*SI Appendix, Fig. S2*), which is expected given the relative higher complexity of both types of nonlinear models. The rather low proportion of the variance explained when fitting the models for each time series is due to both the short length of the time series (especially for the ICES Atlantic Cod dataset with a maximum time series length of 40 y) and their intrinsic variability (*SI Appendix, Fig. S3*). The complete list of all classified time series in each of the three datasets is shown in *SI Appendix, Table S1*.

Dynamical Footprint of Fish Populations. For each time series, the dynamical footprint is composed of 7 different metrics. Those include the dominant eigenvalue [DEV (26)], lag-1 autocorrelation, SD, CV, disparity (27), dispersion, and skewness (*SI Appendix, Table S3*). Metrics were derived both across the entire time series and as a temporal trend. This trend was quantified by calculating Kendall’s tau for each metric, using an overlapping rolling window (window size: 50% of the time series length) to capture how these metrics evolved over time (*Materials and Methods*). In order to meaningfully extract the dynamical footprint from the time series classified as “abrupt,” we filtered out time series with less than 20 measurements before the detected abrupt shift. The temporal distributions of detected




Fig. 2. Dynamical footprint of fish populations. (A) RAM Legacy productivity time series for *Merluccius merluccius* (black line), its best-shape quadratic model (red dashed line) and dynamical footprint. (B and C) ICES Atlantic Cod and FishGlob CPUE time series for *Gadus morhua* and *Rostroraja eglanteria*, respectively (black line), their best-shape abrupt model (red dashed line) and dynamical footprints (bars). In metrics' values calculated using the entire time series correspond to the metrics' percentile ranks across the entire dataset (*Materials and Methods*), and the colors of the bars indicate the percentile ranks within the corresponding shape classes. The values shown for the metric trends corresponds to the Kendall taus (*Materials and Methods*).

abrupt shifts are shown in the supplement, displaying a prevalence of shifts around the 1980s, early 2000s, and around the 2010s, for the RAM, the ICES cod, and the FishGlob datasets, respectively (*SI Appendix, Fig. S4*). Fig. 2A shows an example from the RAM Legacy productivity time series of a stock of European Hake (*Merluccius merluccius*), along with its best-fit quadratic shape (red dashed line) and dynamical footprint (bars). The percentile rank allows for the comparison of each component of the dynamical footprint of a given population to those of all other populations within the same dataset. For instance, although the population exhibits similar trends for dispersion and SD, both metrics fall within different percentile rank bins, which means that the trend in SD is closer to the median SD trend of all populations classified as “abrupt.” Similar examples are presented in Fig. 2B and C, showing the CPUE data from the ICES Atlantic Cod dataset (Sub-area “39F2”) and the FishGlob time series for the Clearnose Skate (*Rostroraja eglanteria*).

Given that some of the metrics composing the dynamical footprint have been used as indicators of abrupt shifts, we checked whether the metrics individually differed across the four shape classes using Mann–Whitney *U* tests. Overall, individual dynamical metrics performed poorly in discriminating among the shape classes across all three datasets (*SI Appendix, Fig. S5*). Metrics calculated using the entire time series tend to show greater differences across classes than metric trends (*SI Appendix, Fig. S5*), although this varies across datasets. More importantly, no single metric or metric trend examined in isolation was able to statistically discriminate all four shape classes (*SI Appendix, Fig. S5*). We observed, however, that certain metrics—notably skewness, disparity, SD, and dispersion—differed significantly between increasing and decreasing abrupt shifts in some cases (*SI Appendix, Fig. S6*).

Machine Learning Model Prediction of Shape Classes from Dynamical Footprints. Although it was not possible to differentiate the four shape classes based on individual dynamical metrics or their trends, we tested whether the entire dynamical footprint could do better in differentiating between shape classes. To evaluate this, we implemented for each dataset, a BRT model that takes the dynamical footprints as input and predicts the four classes as output. We acknowledge that some of the metrics in the dynamical footprint are correlated, and indeed, correlation matrices show commonalities across datasets among related metrics such as CV, SD, and dispersion (*SI Appendix, Fig. S7*). However, PCAs revealed dataset-specific relationships between metrics that could be caused by idiosyncrasies in terms of dimensionality, noise, and/or nonlinear patterns unique to each dataset. We tested the predictive power of the BRT models using leave-one-out cross-validation. For each dataset, stocks were classified as either “abrupt” or “nonabrupt” (the latter encompassing quadratic, linear, and no-change shapes). The results of the classification are presented in Fig. 3, while *SI Appendix, Fig. S8* details the results across all four individual classes. The Receiver Operating Characteristic (ROC) curves in Fig. 3 exhibited moderate performance (AUC of 0.84, 0.75, and 0.65 for the RAM Legacy, the ICES cod, and the FishGlob datasets, respectively), indicating a relatively good (specially for RAM and ICES datasets), but far from absolute, correspondence between the time series abrupt shapes and their dynamical footprints within datasets.

Ideally, if a dynamical system exhibits perfect correspondence between actual shape class and predicted shape class, the BRT model should be able to correctly predict the shape class based solely on the dynamical footprint. To test this, we simulated a fish population dataset using a modified Ricker model capable of




Fig. 3. Machine learning models predict shape classes using dynamical footprints. (A) ROC curves and AUCs are shown for the RAM Legacy (solid), the ICES Atlantic cod (dashed), and the FishGlob datasets (dotted). (B) Confusion matrices presenting aggregated results for two classes: “abrupt” and “nonabrupt.” Potentially abrupt cases are highlighted in light purple (*Bottom Right*).

abrupt shifts (28), and generated one hundred time series of each of the four shape classes, for a variety of model parameters and two different levels of demographic noise (low–high) (*Materials and Methods*). Indeed, the BRT model is able to correctly uncover the relationship between shape and dynamical footprint (*SI Appendix, Fig. S9*, 96% accuracy, $AUC \approx 1$), with an expected performance decrease when high levels of demographic noise are present (75% accuracy, $AUC \approx 0.93$).

In contrast to the simulated results, the matrices shown in Fig. 3B are not conventional confusion matrices. The reason is that in the empirical data we do not know which mechanism is responsible for each shape class. For example, an abrupt shift could result from an abrupt change in the external stressors, say a sudden change in the exploitation rate of a given population, or it could be due to a continuous change in the stressors, say a gradual increase in sea surface temperature that drives the system through a fold bifurcation. Thus, time series classified as “abrupt” but predicted as “nonabrupt” might relate to dynamical changes so sudden that they lie beyond the detection capabilities of the dynamical footprint analysis (*Top Left* cell in Fig. 3). Moreover, time series classified based on shape as nonabrupt but predicted based on their dynamical footprint as abrupt (*Bottom* rows in Fig. 3) could be prone to abrupt shifts that have not yet happened. Thus, “misclassifications” (the off-diagonal elements in Fig. 3) may actually be interpreted as cases where the dynamical footprint does not correspond to the expected shape class of the fish populations.

While our approach primarily focuses on identifying populations most vulnerable to experiencing abrupt shifts, a natural extension is to explore the early prediction of such events, a task inherently challenging due to limited data availability, mainly in terms of time series length. To investigate this predictive capability, we selected populations that were correctly classified as “abrupt” (*Bottom Left* corner in Fig. 3B) and systematically shortened their time series by removing data points backward from the abrupt shift, up to a maximum of 10 y prior. To maintain our minimum length criterion of 20 time points, only original time series exceeding 30 data points were included (19 from RAMLDB, 1 from ICES, and 14 from FishGlob). After each trimming, the dynamical footprint for the shortened series was

recalculated and its classification tested using the corresponding classifier trained on its original dataset, excluding the specific time series under evaluation from this training process. The classification outputs were then recorded for each progressively trimmed length. We found that most populations (24 out of 34) were correctly classified as abrupt even up to 10 y before the actual shift (*Bottom Left* corner in *SI Appendix, Fig. S10*), demonstrating a potential predictive utility of our approach.

Inference of Future Abrupt Shifts via Dynamical Footprint Analysis and Their Spatial Distribution. We can use this “misclassification” to infer the potential for an upcoming abrupt shift. The BRT model generates a probability vector across shape classes for each analyzed population, and these represent the model’s confidence about the populations belonging to each specific shape class (*SI Appendix, Fig. S11 A–C*). Globally (Fig. 4), the distribution of populations classified as “potential abrupt” is widespread across various Large Marine Ecosystems, from the Californian Current to other areas like the Eastern Bering Sea. Many populations from the Northeast and Southeast US Continental Shelf regions were also classified as “abrupt.” The species showing abrupt changes include various rockfish species (e.g., Canary Rockfish, Splitnose Rockfish), as well as other commercially important species such as Hake, Herring, and several species of flatfish. This suggests that abrupt shifts are not limited to specific geographic locations or taxa, but rather a common phenomenon across different marine ecosystems and a wide range of fish populations. For instance, the RAM Legacy productivity time series for the Haddock (*Melanogrammus aeglefinus*), shows no significant change in trend according to the shape classification (*SI Appendix, Fig. S11D*). However, our BRT model classifies this time series as “abrupt” with a probability of around 63%, and given the moderate sensitivity of the model to detect “abrupt” cases (61%), we could consider that this population may actually change abruptly in the future. In other words, this time series is dynamically more similar to those that had a detectable abrupt shift than to the ones that did not display any significant change. Similar examples can be found for the ICES cod and the FishGlob datasets (*SI Appendix, Fig. S11 E and F*). In sum, high classification probability values combined with




Fig. 4. Global distribution of abrupt and potential abrupt time series. (Top) Spatial distribution of time series by Large Marine Ecosystem region. The size of the pie charts is indicative of the number of time series per region N . (Bottom) RAM Legacy productivity time series for (*Melanogrammus aeglefinus*), ICES Atlantic Cod and FishGlob CPUE time series for (*Gadus morhua*) and (*Alosa sapidissima*), with their corresponding best-shape “no change” model (red dashed line). Borders indicate to which dataset the time series belong.

high system-dependent sensitivities of the BRT model increase the likelihood that apparent false positives genuinely represent fish populations prone to abrupt shifts.

Discussion

Dynamical stability has long been central to understanding how ecological populations respond to perturbations (29), and especially to the detection of critical transitions, where small changes in external conditions may trigger abrupt and often irreversible shifts in dynamics (22, 30). Such abrupt shifts have been increasingly observed in natural populations and especially in fisheries (13–15), but our capacity to detect them remains limited (30). Here, we assessed the potential of abrupt shifts in a fish population by classifying its time series based on trajectory shape and comparing its dynamical footprint to those of other populations within the same dataset.

Our assessment refrains from any assumption that the dynamical stability signatures should follow specific patterns, and it remains agnostic about the variety of potential external drivers and disturbances that may influence the dynamics of fish populations. Instead, we only hypothesize that the integration of a diverse set of dynamical metrics, namely the dynamical footprint, will be specific to the dominant shape of a fish population trajectory. In this way, we identified a total of 50 fish populations at risk of future abrupt shifts (13 shifts in productivity in RAM legacy, 6 and 31 shifts in CPUE in ICES and FishGlob, respectively). These include commercially significant stocks from the RAM Legacy database such as the Black Rockfish (*Sebastes melanops*) in the US West Coast, the Haddock (*M. aeglefinus*) in the North Atlantic, and the Pink

cusk-eel (*Genypterus blacodes*) in Australia. The FishGlob dataset also revealed valuable populations at risk, including the Atlantic halibut (*Hippoglossus hippoglossus*) and the Bluefish (*Pomatomus saltatrix*). Notably, the Atlantic cod (*Gadus morhua*) populations were identified as at-risk to abrupt shifts across all three datasets. This consistent appearance across the three different datasets may reflect the often documented sensitivity of Atlantic cod populations to fishing exploitation (11, 31).

While abrupt change in our analysis has the specific mathematical interpretation as a shift in the mean that takes place between two (or few) years, whether such a shift is “abrupt” in the more colloquial meaning of the word depends on the life history of the species. Species with short generation times often experience sudden changes in abundance compared to those with slower life histories which benefit from a demographic buffering to environmental change (32). This distinction highlights a key consideration: Our current classification, focused on quick shifts, might not identify steep, but more progressively unfolding, changes that could nevertheless indicate a significant regime shift. Importantly, our framework can be adapted by integrating alternative classification algorithms more adept at capturing such diverse ecological dynamics (12). Moreover, our current analysis focuses on single abrupt shifts due to the limited length of time series here considered. While adaptable for multiple shifts, our preshift analysis step (Fig. 1) necessitates longer datasets for robust multishift detection, a potential area for future investigations as longer and more comprehensive time series become accessible.

The extent to which our conclusions about the potential risk of an abrupt shift is trustworthy is of course contingent on the model’s sensitivity (i.e., the number of true positives divided by

the number of samples in a given class) and accuracy (i.e., the total number of true positives divided by the total number of samples). Both sensitivity and accuracy are highly influenced by dataset properties, especially noise levels, time series length, and number of populations in each dataset. Despite working with relatively short (*SI Appendix, Fig. S3*), noisy time series, and rather limited dataset sizes (approximately 300 populations in each dataset), the BRT machine learning models produced ROC curves with moderate AUC, reflecting the models' ability to discriminate between positive and negative cases (Fig. 3). To test whether the abrupt shift detection of a given time series is primarily influenced by data limitations—such as short length or high noise levels—we generated a large amount of surrogates (preserving the original time series' mean and variance) and retrieved their predicted shape class. For example, out of ten thousand surrogates of the RAM Legacy productivity time series for the Haddock (*SI Appendix, Fig. S11D*), only 2.8% were also classified as “abrupt,” indicating that the temporal structure (rather than the length or noise to signal ratio) of the time series was relevant in its classification as “abrupt.” Overall, we found no significant correlations between the time series length and the classification probability provided by the BRT models across datasets (*SI Appendix, Fig. S12*). Given the substantial effort required to gather large amounts of fish populations data through global surveys, our approach demonstrates significant value in handling and extracting insights even from constrained datasets.

Yet, the most important factor affecting the performance of the BRT models is probably the heterogeneity of each dataset. This could perhaps explain the lowest sensitivity for the “abrupt” shape class we found in the FishGlob dataset (40%) as this dataset is a sample from more than 8,000 populations indicating a much higher heterogeneity in terms of data-collection methods, species variation, or spatial distribution compared to the other two datasets. One way to improve the performance of the models is to include life-history traits, spatial distribution, and environmental drivers (33) as predictive variables to train the BRT models, in order to capture further differences among the four shape classes. For example, a similar approach was recently implemented to anticipate abrupt shifts in dryland ecosystem functioning through a combination of dynamical metrics and proxies of environmental factors (34). Such additions together with system-specific hyperparameter optimization of the BRT models could maximize accuracy and sensitivity for each dataset and improve model predictions.

Our intention in this study was to show how context-dependent assessments of the patterns underlying fish population dynamics could be used to identify populations prone to future abrupt shifts. Such context dependence is important because relationships among the metrics of the dynamical footprint were not the same across the three empirical datasets (*SI Appendix, Fig. S7*). This result aligned with our a priori expectation that each dataset has particular dynamical idiosyncrasies that are shaped by population specific characteristics, measured variables, and data-collection methods. We therefore encourage future research on a case-by-case dataset-specific dynamical footprint analysis that could also integrate life-history traits, spatial distribution, and environmental data as predictive variables to train machine-learning models. This dataset-specific approach could be adapted to analyze diverse ecological data to enhance the detection of abrupt shifts beyond fish populations, providing a more comprehensive understanding of ecological dynamics that could assist in developing better management and conservation strategies.

Materials and Methods

Data.

RAM legacy database. Following (14), time series data for catch and stock biomass were obtained from the publicly accessible RAM Legacy Stock Assessment Database (RAMLDB v4.61) (23). Then, stock productivity (surplus production) was calculated using the formula: $S_t = B_{t-1} + B_t + C_t$, where B_t represents the total biomass of the population in year t , and C_t represents the catch in year t . The database compiles stock assessments for various taxa, including mollusks, crustaceans, and fish. For consistency and focus, we limited our analysis to fish populations. Out of an initial set of 397 populations, 128 were excluded due to either time series shorter than 20 y (114 populations) or the use of deterministic models to estimate biomass (14 populations), which may introduce biases. The final dataset used included 269 fish populations, and the mean length of the time series was 49 y, with the longest one spanning 141 y. Following the method described in ref. 35, productivity was standardized by average stock biomass.

ICES atlantic cod database. We retrieved time series data for Atlantic cod (*G. morhua*) abundance from the International Bottom Trawl Survey (IBTS), available through the International Council for the Exploration of the Sea (ICES) database (ICES, <https://data.ices.dk>) (24). The data were based on the winter surveys, conducted annually between January and February, spanning the period from 1983 to 2022. These surveys provide a proxy of cod abundance in terms of the number of fish trawled per surveyed rectangle across predefined geographic regions—specifically 1° longitude by 0.5° latitude rectangles within the North Sea. The final dataset used included 136 time series, with a mean length of 38 y and the longest one spanning 40 y.

FishGlob database. The FishGlob database is a freely available resource that provides extensive information on marine fish abundance and biomass, integrating data from various international bottom trawl surveys conducted by different monitoring programs, including those from North America, Europe, and Asia (25). Here, the key variable used was the fish weight per unit of area, which is a standard measure of catch per unit of area swept by the trawl. In cases where fish weight per unit of area was not already computed, it was derived from total fish weight and the area swept during surveys, excluding cases with missing or zero values. First, the data were grouped by survey unit, species, and year, and then the mean fish weight per unit of area for each year was calculated. Thus, each individual time series shows the temporal change of the mean fish weight per unit of area by species and survey unit. The final dataset included 262 fish populations, and the mean length of the time series was 45 y, with the longest one spanning 57 y.

Simulated dataset. We created a simulated library of time series where the actual underlying shape classes were known and were related to specific dynamical changes in the system's stability. The simulations were based on the discrete Ricker-type population model (36), a widely applied framework in population dynamics studies. This model simulates population biomass dynamics driven by logistic growth and incorporates fishing mortality, which allows it to capture various bifurcation types commonly observed in ecological systems (28). We generated 40 deterministic scenarios corresponding to the four distinct shapes targeted for classification. To reflect realistic ecological dynamics we introduced two levels of demographic noise (low and high). Initially, each time series spanned 120 time points, and the breakpoints for the abrupt ones were between the 80th and 100th time point depending on noise, growth, and mortality parameters. Then, to analyze a more homogeneous dataset, for all time series, we selected the first 60 time points. In total, the simulated library was composed of 800 time series with different combinations of noise, growth, and mortality for 10 replicates. Detailed parameter values for the population model can be found in *SI Appendix, Table S2*.

Time Series Shape Classification. We fitted time series with four distinct models: intercept-only, linear, quadratic, and abrupt. To assess the fit of each model, we calculated the Akaike Information Criterion corrected for small sample sizes (AICc) for each fit, with the lowest AICc score indicating the best model for each trajectory. Model validation was performed based on methods described in ref. 14. Specifically, abrupt trajectories were validated by applying an independent breakpoint-detection method (“asdetect”) from ref. 37, ensuring that both the model fit and “asdetect” identified the same shift date within a

tolerance of five years. For cases where the trajectory was classified as quadratic or linear, we evaluated the significance of the higher-order coefficients to confirm the model fit. The “asdetect” method was applied using default parameters—anomalous rate of change set to three median absolute deviations and a detection threshold of 0.15, as in ref. 14. This framework allowed for the identification of at most one shift per time series, with the assumption that any detected shift represents the largest and most significant one. To further assess the reliability of model classifications, we computed two additional classification quality metrics: 1) AICc weight, which quantifies the relative support for each model and 2) normalized RMSE (nRMSE), which represents the ratio of unexplained variance to total variance. The entire classification of the three datasets can be found in *SI Appendix, Table S1*.

Dynamical Footprints. The dynamical footprint was composed of seven different metrics including the DEV, lag-1 autocorrelation, SD, CV, disparity, dispersion, and skewness. Each metric is then calculated both using the entire time series, and in a time-varying way using a rolling window approach. Metrics such as DEV, dispersion, CV, and disparity were calculated using the raw productivity time series, whereas the lag-1 autocorrelation, SD, and skewness were calculated from the residuals after detrending. The DEV was computed from the estimated Jacobian using the Smap function as in refs. 38 and 39. For this, the optimal embedding dimension, E , and time lag, τ and nonlinearity term θ , were estimated using the entire time series. To determine these hyperparameters, S-map prediction skill was evaluated for combinations of embedding dimensions E (ranging from 1 to 6), time lags τ (1 and 2), and nonlinearity term θ (from 0 to 2.5). The set of E , τ , θ giving the best S-map forecast skill was chosen. After retrieving the time series of Jacobian, we calculated the dominant eigenvalue at each time point. Then, we averaged all eigenvalues derived in the time series as our DEV metric. The disparity index (27) was computed using the function from the “lightdosimetry” R package. The output in the dynamical footprint for

the metrics calculated using the entire time series corresponds to the percentile rank across the entire dataset, whereas the value of the metrics calculated using a rolling window corresponds to the Kendall tau across the second half of the time series.

Machine Learning Model. To predict the shape classes for each dataset we used a BRT model, employing a leave-one-out approach for cross-validation. Initially, the dataset was prepared by selecting the metrics in the dynamical footprint as predictors and the shape class as target variable. The BRT model was trained using the “train” function from the “caret” package, applying the “xgbTree” method to leverage the boosting capabilities of the “XGBoost” library. The hyperparameters for the model were set using a predefined grid (100 boosting rounds, maximum tree depth = 3, learning rate = 0.1, and minimum loss reduction = 0). Following model training, we calculated the classification probabilities for the left-out observation using the trained model repeatedly for each observation in the dataset. Then, we calculated the maximum classification probability per sample and determined the shape class based on the highest one.

Data, Materials, and Software Availability. All the computational analyses and simulations were performed in R 4.4.1. The fully reproducible workflow, including code and processed data, is publicly available at GitHub, <https://github.com/alejvcano/dynfoot2025> (40). All other data are included in the manuscript and/or *SI Appendix*.

ACKNOWLEDGMENTS. We thank Benoît Pichon, Hsiao Hang Tao, and Mathieu Péliissié for providing feedback and the two anonymous reviewers whose comments greatly improved the manuscript. This study was financially supported by J.C.J. ANR-22-CE32-0001-01.

1. R. Hilborn *et al.*, Effective fisheries management instrumental in improving fish stock status. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 2218–2224 (2020).
2. M. L. D. Palomares *et al.*, Fishery biomass trends of exploited fish populations in marine ecoregions, climatic zones and ocean basins. *Estuar. Coast. Shelf Sci.* **243**, 106896 (2020).
3. Food, AOot Nations, *The State of World Fisheries and Aquaculture 2024: Blue Transformation in Action* (FAO, Rome, 2024).
4. M. C. Melnychuk *et al.*, Identifying management actions that promote sustainable fisheries. *Nat. Sustain.* **4**, 440–449 (2021).
5. E. Ojea, S. E. Lester, D. Salgueiro-Otero, Adaptation of fishing communities to climate-driven shifts in target species. *One Earth* **2**, 544–556 (2020).
6. B. S. Halpern *et al.*, A global map of human impact on marine ecosystems. *Science* **319**, 948–952 (2008).
7. S. Georgian *et al.*, Scientists’ warning of an imperiled ocean. *Biol. Conserv.* **272**, 109595 (2022).
8. C. M. Free *et al.*, Impacts of historical warming on marine fisheries production. *Science* **363**, 979–983 (2019).
9. K. Brander, Impacts of climate change on fisheries. *J. Mar. Syst.* **79**, 389–402 (2010).
10. A. J. Pershing *et al.*, Slow adaptation in the face of rapid warming leads to collapse of the Gulf of Maine cod fishery. *Science* **350**, 809–812 (2015).
11. C. Sguotti *et al.*, Catastrophic dynamics limit Atlantic cod recovery. *Proc. R. Soc. B Biol. Sci.* **286**, 20182877 (2019).
12. C. Sguotti, L. Färber, G. Romagnoni, “Regime shifts in coastal marine ecosystems: Theory, methods and management perspectives” in *Treatise on Estuarine and Coastal Science (Second Edition)*, D. Baird, M. Elliott, Eds. (Elsevier, 2024), pp. 50–72.
13. K. A. Vert-pre, R. O. Amoroso, O. P. Jensen, R. Hilborn, Frequency and intensity of productivity regime shifts in marine fish stocks. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 1779–1784 (2013).
14. M. Péliissié, V. Devictor, V. Dakos, A systematic approach for detecting abrupt shifts in ecological timeseries. *Biol. Conserv.* **290**, 110429 (2024).
15. C. Möllmann *et al.*, Tipping point realized in cod fishery. *Sci. Rep.* **11**, 14259 (2021).
16. G. Beaugrand, The North Sea regime shift: Evidence, causes, mechanisms and consequences. *Prog. Oceanogr.* **60**, 245–262 (2004).
17. R. Hilborn, C. J. Walters, *Quantitative Fisheries Stock Assessment: Choice, Dynamics and Uncertainty* (Springer Science & Business Media, 2003).
18. C. F. Clements, J. L. Blanchard, K. L. Nash, M. A. Hindell, A. Ozgul, Body size shifts and early warning signals precede the historic collapse of whale stocks. *Nat. Ecol. Evol.* **1**, 1–6 (2017).
19. M. Krkošek, J. M. Drake, On signals of phase transitions in salmon population dynamics. *Proc. R. Soc. B Biol. Sci.* **281**, 20133221 (2014).
20. C. F. Clements, M. A. McCarthy, J. L. Blanchard, Early warning signals of recovery in complex systems. *Nat. Commun.* **10**, 1681 (2019).
21. F. Zhang, Early warning signals of population productivity regime shifts in global fisheries. *Ecol. Indic.* **115**, 106371 (2020).
22. M. Scheffer *et al.*, Early-warning signals for critical transitions. *Nature* **461**, 53–59 (2009).
23. 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 (2012), <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-2979.2011.00435.x>.
24. ICES, *International Bottom Trawl Survey Working Group (IBTSWG)* (ICES Scientific Reports, 2020).
25. A. A. Maureaud *et al.*, FISHGLOB_data: An integrated dataset of fish biodiversity sampled with scientific bottom-trawl surveys. *Sci. Data* **11**, 24 (2024).
26. F. Grziwotz *et al.*, Anticipating the occurrence and type of critical transitions. *Sci. Adv.* **9**, eabq4558 (2023).
27. M. Fernández-Martínez, J. Peñuelas, Measuring temporal patterns in ecology: The case of mast seeding. *Ecol. Evol.* **11**, 2990–2996 (2021), <https://onlinelibrary.wiley.com/doi/pdf/10.1002/ece3.7291>.
28. V. Dakos, S. M. Glaser, C. Hsieh, G. Sugihara, Elevated nonlinearity as an indicator of shifts in the dynamics of populations under stress. *J. R. Soc. Interface* **14**, 20160845 (2017).
29. S. L. Pimm, The complexity and stability of ecosystems. *Nature* **307**, 321–326 (1984).
30. V. Dakos *et al.*, Tipping point detection and early warnings in climate, ecological, and human systems. *Earth Syst. Dyn.* **15**, 1117–1135 (2024).
31. K. T. Frank, B. Petrie, W. C. Leggett, D. G. Boyce, Large scale, synchronous variability of marine fish populations driven by commercial exploitation. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 8248–8253 (2016).
32. W. F. Morris *et al.*, Longevity can buffer plant and animal populations against changing climatic variability. *Ecology* **89**, 19–25 (2008), <https://onlinelibrary.wiley.com/doi/pdf/10.1890/07-0774.1>.
33. M. Péliissié, V. Devictor, O. Jensen, V. Dakos, Ocean warming drives abrupt declines in fish productivity at global scale. *EcoEvoRxiv* [Preprint] (2025), <https://doi.org/10.32942/X22586> (Accessed 28 May 2025).
34. P. N. Bernardino *et al.*, Predictability of abrupt shifts in dryland ecosystem functioning. *Nat. Clim. Chang.* **15**, 86–91 (2025).
35. T. E. Essington *et al.*, Fishing amplifies forage fish population collapses. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 6648–6652 (2015).
36. W. E. Ricker, Stock and recruitment. *J. Fish. Res. Board Can.* **11**, 559–623 (1954).
37. C. A. Boulton, T. M. Lenton, A new method for detecting abrupt shifts in time series. *F1000Res.* **8**, 746 (2019).
38. G. Sugihara, B. T. Grenfell, R. M. May, H. Tong, Nonlinear forecasting for the classification of natural time series. *Philos. Trans. R. Soc. Lond. Ser. A Phys. Eng. Sci.* **348**, 477–495 (1997).
39. D. A. O’Brien *et al.*, EWSmethods: An R package to forecast tipping points at the community level using early warning signals, resilience measures, and machine learning models. *Ecography* **2023**, e06674 (2023), <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ecog.06674>.
40. A. V. Cano, dynfoot2025. Github. <https://github.com/alejvcano/dynfoot2025>. Deposited 29 July 2025.