complexities

Perspective

Resilience, Tipping Points, and Hysteresis

Peter Grindrod [1,2]

[1] Astut Ltd., Oxford OX10 6GF, UK; peter.grindrod@maths.ox.ac.uk
[2] Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

Abstract

In the essay we introduce present-day systems concepts, such as resilience, tipping points, and hysteresis effects, via the concept of fast–slow dynamical systems (whether explicit in the models or implicit through bifurcation and stability behaviours). These lead naturally to ideas first propagated within catastrophe theory, fifty years ago. We discuss the historical *catastrophe* (the backlash) that befell such an abstract yet mathematically grounded (and thus inescapable) theory within economics and also its subsequent re-appraisal and re-adoption. Finally, we discuss some of the challenges inherent in anticipating tipping points from live systems data (observations), within systems-theoretic interpretations, and whether methods from topological data analysis might respond to them. While it is fashionable for national, governmental and policy institutions to speak of "resilience" in all manner of national systems contexts, we aver that it is foolishly inadequate to do so without an understanding and consideration of tipping points and hysteresis (sometimes termed "path dependence"), giving rise to "lock-in".

Keywords: resilience; tipping points; hysteresis

1. Introduction

Resilience is the capacity of a system or entity to withstand or to recover (if ever and how quickly?) from insults and sudden changes; it is a measure of its toughness or resistance to change as a result of externalities and impacts. Sometimes it is defined as the ability of the entity to restore itself to normal after such disruptions, whose occurrences represent real risks (perturbations and shocks to the entity), and is seen as a kind of elasticity. Yet of course, physical elasticity is most often linear and results almost inevitably in directly reversible phenomena. Many systems in organisational, social, and economic fields have no immediate way to restore themselves.

The latter idea of possibly reversible recovery does not necessarily work within many nonlinear systems, since there may be some *hysteresis* present. It is a simple category error to assume that any system is reversible within the short term, exploiting restorative dynamics, and that any recovery might be as immediate as the initial disruption was.

In general, we say that a parameterised dynamical system exhibits hysteresis if there exists a parameter interval on which multiple stable invariant sets coexist and the limiting state selected by solution trajectories depends on the direction of parameter variation, so that quasi-static forward and backward parameter sweeps generate distinct solution trajectories.

The popular term "tipping point" invites a mental image of a seesaw with the fulcrum perhaps moving slowly or with extra weights being added to one or other end of the seesaw. At some point, what was up in the air must suddenly tip down. But this is easily reversible:

Check for updates

Academic Editor: Jarosław Kwapień

Received: 12 February 2026
Revised: 20 March 2026
Accepted: 31 March 2026
Published: 3 April 2026

you just reduce the weight added or you move the fulcrum back a bit and it tips right back the other way. We will show that this simplicity is not natural in nonlinear systems, where slow dynamics may result in sudden fast "catastrophes", but there is no simple reversibility and no easy return path. Instead, one must navigate within the *new normal*, until you find a way back, usually via a reverse catastrophe located somewhere else within the full state space.

2. Fast and Slow Processes Beget Catastrophes

One way to imagine nonlinear situations is to think of the existence of both fast and slow active processes within the systems, whether socio-economic, biological/chemical, financial, emotional/psychological, or physical. Fast processes are effectively instantaneous dynamics (perhaps relative to the time scales of some such natural processes or to any observer's ability to measure them or to become aware of fast state changes), and they must be an order of magnitude or more faster than the slow processes (which, perhaps, operate on time scales that are observable). *Fast* and *slow* are thus relative terms. Alternatively, the fast–slow partition of processes may arise via an asymptomatic limit of some small system parameters controlling a subset of the process rates of change.

Within any high-dimensional system state space, say a Euclidean space, a *manifold* is a smooth lower dimensional subset in which (everywhere) locally resembles a lower dimensional Euclidean space. Think of a smooth curved two-dimensional surface, such as a tablecloth, that is smoothly curved and is embedded as a two-dimensional surface within a three-dimensional space.

Now, whenever we observe, we will find that the fast system processes are usually in equilibrium, meaning that the state of the system must remain on a manifold, M, which is defined as the set of states where all of those fast processes are at equilibrium. Meanwhile, the slow processes will continually move the state around on M, whether due to internal (autonomous) dynamics or to external forcing terms.

The 1970s saw the development of the mathematical science of catastrophe theory [1]. It is a branch of bifurcation theory within dynamical systems theory and is a special case of singularity theory within geometry. It generally considers cases where stable equilibrium minimises a smooth, well-defined potential or "Lyapunov" function. But that is not necessary: only the fast–slow split matters, whether made explicit, via the system equations, or implicit, via *centre manifold* theoretical considerations.

Within the fast–slow perspective given here, catastrophes occur because the equilibrating fast evolution of a system may result in folds of the fast-equilibrium manifold, M, along with fast dynamics off it running tangential to M at some points along those folds. This results in the state of the systems *falling off* M at such folds and then almost instantaneously falling back onto M elsewhere, perhaps a long way away, after a very fast transition. Once occupying the new location, the system state evolves again due to the slow dynamics and may move around until it falls back at some other point, on some other fold. It is clear that the catastrophe, the very fast transition, is not directly reversible at all and that the folds are most likely to be discrete, with the state subject to slow dynamical movements between them. Usefully, catastrophe theory also provides a rigorous taxonomy of different types of folds within manifolds of increasing complexity that might occur.

Most mathematics undergraduate courses today will contain an element of fast–slow process splitting within dynamical systems lectures that suggests multiple time-scale asymptotic approaches or that occurs when systems are close to centre manifolds on which slow processes evolve close to equilibria (part of bifurcation theory).

In Section 3 we discuss how catastrophe theory fell into bad odour within its application to economics, in particular. In part, this was due to resistance from the economics community to the use of dynamical systems theory to remove the comfort blanket provided by small perturbation theory (and linear stability) and the Central Limit Theory-generated *Gaussian paradise*. In part, it was due to its rather abstract approach and the consequent need for assumptions about fast–slow partitions within active processes and/or the mathematicians' desire to assume the existence of extant fast equilibrating "potentials" which did not coincide with observable qualities or existing economic theoretical concepts. As a result, the economics academic mainstream wriggled free from the imposition of catastrophe theory and entered its own long dark tunnel of generalised linear modelling (suitable perhaps for small disruptions) and its tedious application within econometrics. In fact, the application of econometrics (statistical methods addressing economic data) has itself seen a rise and fall in its popularity, with periods of scepticism.

Nevertheless, when stripped down to simple facts of life for fast–slow nonlinear dynamical systems, the ideas behind catastrophe remained valid and are inevitable.

Over the past two decades, this rejection in economics has been revised, especially since the criticism and ridicule set out in [2], and there is now some excitement [3], more recently in [4], across a wide range of other applications to complex systems.

Nowadays, it is common to talk of *tipping points* [5], but this simplicity might ignore one of the key consequences of catastrophes: hysteresis.

The terms *resilience* (the restorative ability to repair oneself after insults) and *tipping points* really need to embrace hysteresis. Moreover, the catastrophe theoretic approach (the nature of the occurrence of folds in the fast-equilibrium manifold) implies that a tipping point can occur very naturally from within fast–slow systems as well as be driven by external forcing.

The general mathematical setting is rather straightforward. We have some time-dependent state variables, say $\mathbf{x}(t) \in \mathbb{R}^n$, which evolve on a fast time scale, and some other time-dependent state variables, $\mathbf{y}(t) \in \mathbb{R}^m$, which evolve on a slow time scale. Some of each may be unobservable. Then, for a small real parameter, $\epsilon > 0$, we consider a coupled system of the form

$$\epsilon\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{y}), \quad \dot{\mathbf{y}} = g(t, \mathbf{x}, \mathbf{y}),$$

where $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is smooth and is defined over state space, the $(\mathbf{x}, \mathbf{y})$ space. Within the limit of positive $\epsilon \to 0$, the $\mathbf{x}$ equation is a fast dynamic, and we must have $f(\mathbf{x}, \mathbf{y}) = 0$ at all times, a condition which defines the manifold, M, and its pseudo-equilibrium states, the m dimensional null set of f. M is sometimes called the "nullcline" for f and may be folded and be multivalued as a function of $\mathbf{x}$ for some fixed values of $\mathbf{y}$, with multiple branches, so typically there may be regions for $\mathbf{y}$-values where there is multiple stability, with the $\mathbf{x}$ dynamic moving the system almost instantaneously onto such branches (sub-regions) of M.

Meanwhile, the possibly non-autonomous $\mathbf{y}$ dynamic, given by $g : \mathbb{R} \times \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^m$, is slow and may be subject to external forcing and moves the state around on M, causing it to fall off folds as demanded.

For example, suppose $n = m = 1$ and $f(x, y) = x - x^3 - y$ and $g(x, y) = x$. The one dimensional manifold M is the nullcline of f (see Figure 1). There are two branches of M which are stable for the fast x-dynamic (and one unstable middle branch) for a range of y values, where $y \in (-2/3\sqrt{3}, 2/3\sqrt{3})$. For other values of y, there is a single stable branch of M. Meanwhile, the possibly non-autonomous y-dynamic, given by g, causes the state to wander over M, and occasionally we may see a catastrophe where it falls off at the old fold edge, when $y = \pm 2/3\sqrt{3}$, and switches almost instantaneously to the other available stable branch, as y moves beyond the edge. This is indeed a tipping point and is irreversible.

See the autonomous example in Figure 1. After a transient period, all orbits tend towards a limit cycle with two catastrophes in each period. For a more general choice of $g(t, x, y)$ though, the orbit would wander about on either of the two stable outer branches and fall off the folds onto the other stable branch as and when required.

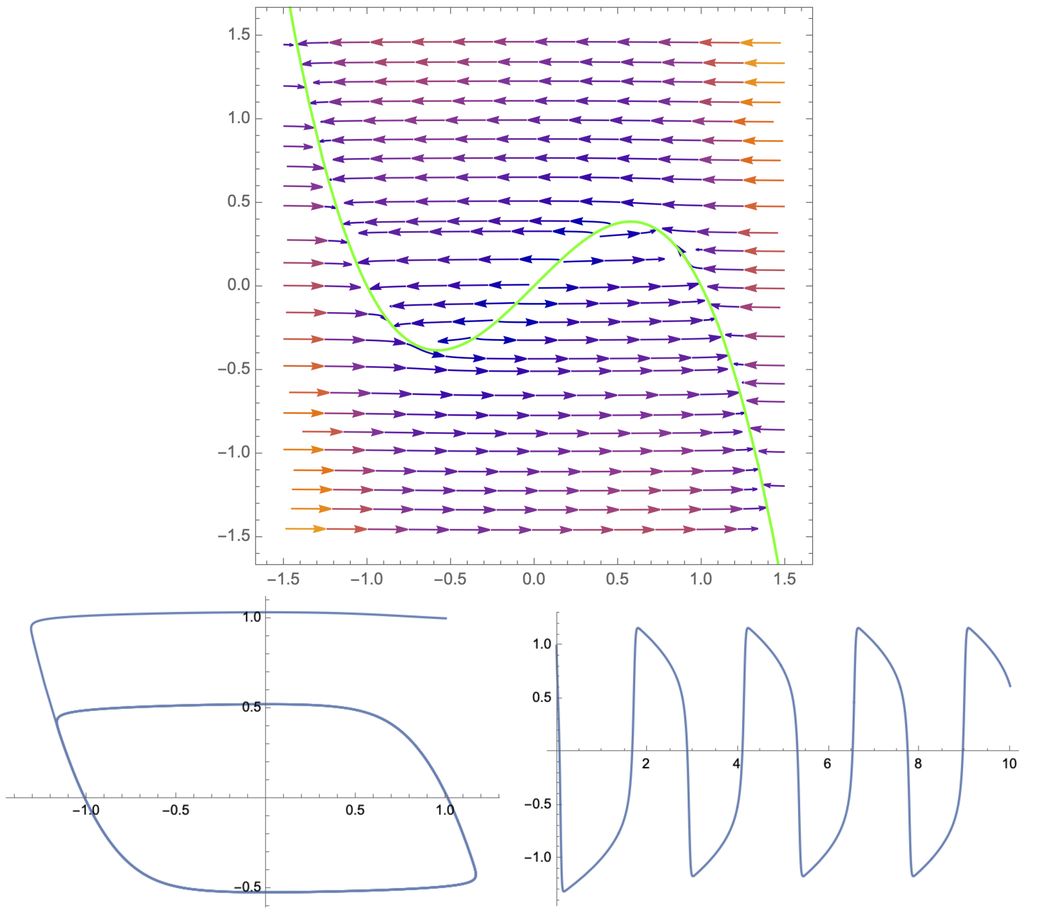

Figure 1. (**Above**): The nullcline for $f(x,y) = x - x^3 - y$, which is the manifold M in this case, is shown in green. Assuming $g(x,y) = x$ (see text) and $\epsilon = 0.05$, we show the resulting vector field. Catastrophes occur each time the orbit reaches either of the folds in M and switches to the other stable branch. (**Below**): An example orbit converging to a (relaxation) limit cycle and the corresponding $x(t)$ coordinate showing the pair of catastrophes within each period.

There is a classification of various distinct types of catastrophe-generating manifolds. In Figure 2 we show one such manifold that contains three distinct surfaces that are stable for the fast dynamic, on which the slower state dynamic can move the state along orbits, resulting in path dependence (hysteresis).

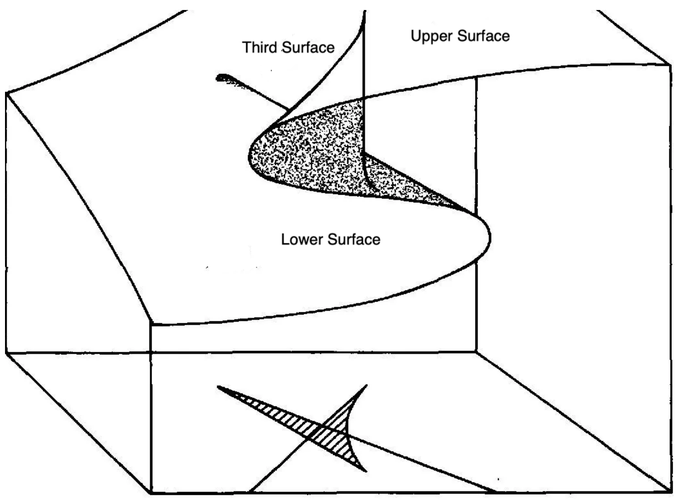

Figure 2. The butterfly surface in a three-dimensional space with three quasi-stable surfaces with possible fast transitions between them.

As described above, it is not possible for a dynamic to return back to the cliff edge from whence it fell; instead, the state must move slowly along its new stable branch until it falls off the other cliff edge back onto its original branch. This is hysteresis.

3. Historical Context

The ideas of catastrophe, tipping points, and hysteresis discussed in this paper have an interesting intellectual history, particularly within economics. In the 1970s, catastrophe theory—developed by René Thom and popularised by Christopher Zeeman—was widely promoted as a way of explaining sudden changes in social and economic systems. It offered a mathematical framework in which gradual changes in underlying conditions could produce abrupt transitions between equilibria and in which recovery might not follow the same path as collapse. These features made it attractive for modelling stock-market crashes, unemployment surges, and shifts in economic regimes.

However, by the early 1980s, catastrophe theory had largely fallen out of favour within mainstream economics. Part of the resistance came from methodological concerns: many early applications were seen as metaphorical rather than empirical, relying on elegant geometry but weakly grounded in data. More importantly, the discipline itself was moving toward equilibrium models built on rational expectations, representative agents, and smooth adjustment. In this intellectual climate, the nonlinear, multi-equilibrium structures of catastrophe theory appeared both analytically inconvenient and philosophically suspect. As a result, the economics mainstream held to linear econometric modelling, a shift that left little room for sudden discontinuities or irreversible transitions.

Despite this rejection, the core ideas of catastrophe theory never disappeared. Instead, they re-emerged under different names. In the late 1980s, Brian Arthur's work [6–8] on increasing returns and path dependence showed how small historical events could lock economies into particular outcomes and how reversing the original conditions might not restore the earlier state. Around the same time, macroeconomists introduced the concept of hysteresis in unemployment, recognising that temporary shocks could have long-lasting or permanent effects. These developments echoed the same underlying structures that catastrophe theory had emphasised: multiple equilibria, irreversible transitions, and strong dependence on history.

Arthur's work can be seen as a reframing of the discontinuities emphasised by catastrophe theory into a more concrete economic language of lock-in and increasing returns. Catastrophe theory focused on sudden jumps between equilibria, often represented geometrically as a trajectory falling off a fold in a manifold. Arthur, by contrast, emphasised gradual reinforcement processes that lead to lock-in. In his models, small early differences are amplified through positive feedback—such as network effects, learning curves, or coordination benefits—until one outcome dominates and becomes effectively irreversible. Instead of a sudden jump caused by crossing a geometric threshold, Arthur's story is one in which positive feedback progressively magnifies small advantages. Structurally, however, the two perspectives are closely related. In catastrophe theory, a system may have two equilibria; as a parameter changes slowly, it eventually crosses a threshold and jumps to the other state. In Arthur's framing, two technologies compete, early random adoptions occur, and positive feedback amplifies one of them until the system locks into a single standard. In both cases, the outcome depends on history, reversal is difficult, and multiple equilibria exist.

The reception of these ideas in economics differed sharply. Catastrophe theory was often viewed as abstract, overhyped, and lacking empirical grounding. Arthur's work, in contrast, was mechanism-based, intuitive, and directly connected to observable phenomena such as technological competition and network effects. His models were also mathematically simpler and easier to interpret. As a result, Arthur succeeded where catastrophe theory had largely failed: he reintroduced nonlinear, path-dependent dynamics into economics without relying on the controversial geometric formalism that had earlier alienated the discipline.

In the decades that followed, empirical realities increasingly forced economists and other scientists to confront such nonlinear behaviour. Financial crises, ecological collapses, and climate tipping points all exhibited sudden regime shifts and slow, path-dependent recovery. Complexity economics, agent-based modelling, and network approaches began to treat economies explicitly as dynamical systems with feedback, instability, and multiple attractors. Under this new paradigm, the once-discredited concepts of tipping points and hysteresis became central rather than peripheral.

The perspective developed here helps clarify why those ideas are unavoidable. By focusing on systems with interacting fast and slow processes, the occurrence of catastrophes becomes a natural consequence of the geometry of the fast-equilibrium manifold. When slow dynamics push the system toward a fold, a rapid transition becomes inevitable, and the return path lies elsewhere in the state space. As emphasised here, the resulting behaviour is not simply reversible: the system must travel along a different trajectory before it can switch back, if it does at all.

In his later writings, Ian Stewart, one of the principal mathematical expositors of catastrophe theory in the 1970s [1], offered a reflective account of the rise and decline of catastrophe theory [9,10], treating it as an instructive episode in the interaction between mathematics, scientific fashion, and application. He described how the theory moved rapidly from its mathematical origins into a wide range of social and behavioural sciences, including psychology, sociology, and economics. This expansion, he suggested, was driven as much by enthusiasm and intellectual fashion as by careful empirical grounding.

In Stewart's account, the underlying mathematics of catastrophe theory remained sound, and some applications were reasonable. The difficulty arose because the theory was often presented as a universal explanatory tool, capable of describing everything from stock-market crashes to social revolutions. Many of these applications relied on loose analogies rather than measurable variables or testable models. As a result, catastrophe theory acquired a reputation for hype and overextension. The subsequent backlash, in Stewart's view, was therefore not a straightforward scientific refutation but partly a sociological reaction against exaggerated claims and poorly grounded applications.

A concise version of this interpretation appears in [9], where Stewart recounts the excitement surrounding catastrophe theory, its rapid diffusion into the social sciences, and its eventual loss of credibility. He presents the episode as a cautionary tale: a mathematically sound theory can suffer reputational damage if it is promoted too broadly or applied without sufficient empirical discipline. The decline of catastrophe theory, in this telling, was less a rejection of its mathematics than a reaction to the way it had been used and publicised.

Seen in this historical light, the modern emphasis on tipping points and resilience does not represent a new discovery so much as a rediscovery. Catastrophe theory was once set aside because it seemed abstract, difficult to test, and inconsistent with prevailing equilibrium models. But as empirical evidence accumulated and computational methods improved, the same core insights—nonlinearity, multiple equilibria, and hysteresis—returned under the language of complexity, resilience, and tipping points. The fast–slow dynamical framework described here provides a clear and practical way to understand those phenomena without the excesses that once surrounded catastrophe theory's earlier applications.

4. Indications of Catastrophes

With this model in mind, notice first that the shocks (catastrophic switches) take place in the fast variables, while the slow variables just keep trundling on.

Suppose that we are on M and the slow dynamic (whether autonomous or externally driven) takes us along a path towards the catastrophic fold/edge. We would observe small movements towards the fold, where the fast forces must be tangential to the manifold, and these would result in an acceleration of the fast variable into the very catastrophe. In time-series terms, we would observe a sharp increase or sharp decrease in the second difference, foretelling an imminent catastrophic positive or negative change respectively.

Consider the time series for a blue chip company's closing share price on 6377 successive trading days, shown (left) in Figure 3. On the right, we plot the forward two-day movement against the backward second difference (the discrete acceleration). The two are positively correlated overall and with relatively few extremes within the top-right and bottom-left quadrants.

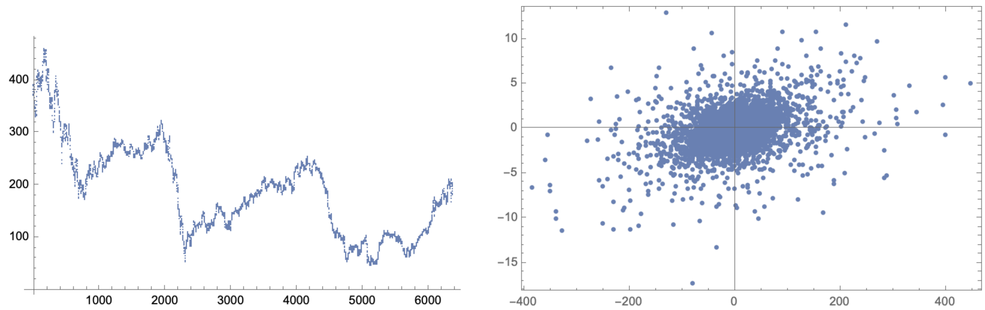

Figure 3. (**Left**): Time series for a blue chip company's closing share price on 6377 successive trading days. (**Right**): We plot the forward two-day movement against the backward two-day second difference (the discrete acceleration). The two are positively correlated overall, with few relative extreme shifts (possible catastrophes) mostly within the top-right and bottom-left quadrants.

5. TDA Methods: Alterative Approaches

Topological Data Analysis (TDA) identifies abrupt changes and catastrophes in (single or multivariate) time series by applying persistent homology (PH) to dynamically embedded time-series windows to track the evolution of topological features like holes and connected components.

At present, TDA (PH) is itself going through its own decadal evolution. While promising theoretically, it provides a concept which can augment and challenge the existing approaches in many fields, especially in dealing with high-dimensional and complex point clouds that may have many distinctive features [11]. It is not always easily accessible or computable (though many variants of algorithms and codes are available) and must prove its worth in applications.

For an observed dynamical system, time-series data is first transformed into a point cloud, usually via a sliding window of a fixed length, e.g., K, creating a set of points that capture the system's dynamics at each time step. Of course, these points largely represent an embedding in $\mathbb{R}^K$ of the stable portions of any attractor manifold (for the fast system). This is consistent with, and justified by, Takens' well-known embedding theorem.

A PH filtration may then be applied to these point clouds, which is a sequence of increasingly larger shapes within the embedding space; the shapes are the "birth" and "death" of topological features, such as connected components (zero-dimensional holes), or loops (one-dimensional holes), or voids (two-dimensional holes), at different scales within this filtration.

Then, any abrupt tipping points (catastrophes) can be detected by comparing the persistence diagrams or other topological features from consecutive time windows.

A useful approach here might be to keep the point cloud within a rather high K-dimensional embedding space and auto-identify and separate the large structures (characterising the topology of any attracting manifolds for the fast system, which could also be high-dimensional) and the small-scale noise; the division between these is a matter of decadal interest, for general time-series analysis.

TDA emerged in the early 2000s as a framework for extracting robust geometric and topological features from data. The foundational idea is that many datasets, especially those generated by dynamical systems, lie near low-dimensional manifolds embedded in higher-dimensional spaces. Persistent homology (PH), the main computational tool in TDA, identifies features such as connected components, loops, and higher-dimensional holes that persist across scales and thus represent structural properties of the underlying system rather than noise.

The mathematical foundations of PH were established in the early 2000s. The seminal paper [12] introduced the concept of PH as a method for computing topological features across multiple scales, while [13] provided an algebraic formulation and efficient algorithms. These works established PH as a practical computational tool rather than a purely theoretical construct.

Soon afterwards, researchers began applying TDA to time series and dynamical systems. The key conceptual step was the recognition that delay-coordinate embeddings, following Takens' theorem, allow time series to be converted into point clouds whose topology reflects the underlying attractor, with the idea that topological invariants of the delay reconstructions might be used to characterise distinct dynamical regimes. Of course, that is possibly more straightforward than detecting the precursors of any sudden tipping points.

One of the earliest systematic treatments of topology for dynamical data was given in 2009 by [14], which presented TDA as a general framework for analysing complex datasets, including those arising from dynamical processes. Around the same time, researchers began applying persistent homology to detect qualitative changes in system behaviour.

A major step toward practical applications came with the work of [15], which applied TDA to financial time series. This introduced the sliding-window embedding approach and tracked changes in persistence landscapes over time. Their results showed that large changes in topological features could precede market crashes, suggesting that PH could provide early-warning signals of systemic transitions. This work demonstrated that TDA could detect regime shifts in noisy, real-world economic data.

In parallel, other researchers began using TDA to detect critical transitions and tipping points in ecological and climate systems. For example, ref. [16] showed that topological features extracted from time series could be used to detect changes in system dynamics before bifurcations.

Similarly, the work in [17] used sliding-window embeddings and persistent homology to identify periodicity and structural changes in time-series data, demonstrating that PH could recover dynamical features such as cycles without explicit modelling.

More recently, TDA has been integrated with early-warning-signal research in complex systems. Studies such as [4] have examined interacting tipping elements in climate and ecological systems, showing how cascading transitions can emerge. Although not always explicitly framed in topological terms, such work provides the dynamical context in which PH-based early-warning indicators become relevant.

Across these studies, a consistent theme emerges: PH can track global structural changes in the geometry of observed dynamics. Instead of relying solely on local statistical

https://doi.org/10.3390/complexities2020010

indicators such as variance or autocorrelation, TDA monitors the evolving topology of reconstructed attractors. This makes it particularly well suited to systems with fast–slow dynamics and catastrophic transitions, where the underlying manifold structure changes as the system approaches a fold or bifurcation. However, the scarcity of the tipping points within time series means that such defining phenomena may be crowded out by the quasi-stable dynamical paradigms between them.

Over the last decade, the use of TDA and PH for detecting tipping points has evolved from a mathematical concept into a structured and empirically grounded research direction.

Rather than treating topological signals as purely empirical markers, recent work has attempted to show why certain topological changes should occur as systems approach critical transitions. For example, in financial bubble models [18] or other nonlinear growth processes, the deformation of the attractor or trajectory geometry near a singularity can produce characteristic changes in the persistence landscape. This type of mechanistic grounding addresses a long-standing criticism of early-warning indicators. In [19], a strictly causal early-warning system for financial crises is proposed, based on topological features extracted from multivariate financial time series. The goal is to show that PH can produce interpretable early-warning signals that outperform simple volatility-based indicators.

At the same time though, the broader literature on tipping points has become more cautious about generic early-warning signals. Rather than presenting PH as a universal predictor, recent work treats it as one component in a multi-indicator framework, valued for its ability to capture global geometric changes in multivariate data.

6. Open Challenges

The fast–slow dynamical systems perspective presented here highlights a number of conceptual and practical challenges for the application of Topological Data Analysis (TDA) and persistent homology (PH) to systems exhibiting tipping points, catastrophes, and hysteresis.

First, in fast–slow systems, catastrophic transitions are typically abrupt, irreversible on short time scales, and separated by long periods of slow evolution. This temporal asymmetry creates a fundamental difficulty for TDA methods based on sliding-window embeddings: most windows capture only quasi-equilibrium motion along stable manifolds, while the catastrophic events occupy a very small fraction of the data. As a result, the topological features extracted from delay embeddings will be dominated by the geometry of the stable attractor rather than by the structures associated with imminent transitions.

Second, the topology observed in delay-coordinate embeddings reflects the structure of the attracting manifold of the fast dynamics, rather than the full state space in which both folds and catastrophes occur. Since tipping points arise from geometric features of the fast-equilibrium manifold, such as folds or bifurcation structures, these critical features may be poorly represented or even invisible in reconstructed attractors, particularly when only partial observations of the system are available.

Third, the presence of hysteresis is subordinate to Arthur's "path dependence" and multi-stability. Systems may move slowly along one stable branch after a catastrophe and only return to the original regime via a different fold located elsewhere in the state space. From a TDA perspective, this means that similar topological signatures in time-series windows may correspond to qualitatively different dynamical contexts, depending on the system's history and its position relative to folds in the manifold. Consequently, interpreting changes in persistence diagrams as early-warning signals requires additional dynamical context.

Fourth, the separation between a large-scale topological structure and small-scale noise remains a practical challenge. In sliding-window embeddings, the point clouds typically represent noisy samples from high-dimensional attractor manifolds. Distinguishing persistent features associated with genuine manifold geometry from those produced by noise or finite sampling is still an open methodological problem.

Fifth, there is a broader conceptual challenge concerning the interpretation of topological indicators, especially within high-dimensional state spaces. While PH can detect structural changes in reconstructed dynamics, the connection between specific topological features and underlying dynamical mechanisms—such as approaching folds, bifurcations, or cascades of tipping elements—remains an active area of research.

Recent work has attempted to provide mechanistic explanations for topological early-warning signals, but the field has moved toward viewing PH as one component within a multi-indicator framework rather than a universal predictor. Can we be more assertive for both low- and high-dimensional state spaces and their attracting sets?

Perhaps Stewart's cautionary take about catastrophe theory should be borne in mind, even as some promising applications of TDA (as PH) are championed in addressing reliance in general and tipping point dynamics in particular.

Funding: This research was partly funded by UKRI/EPSRC grant number EP/Y028872/1 Mathematical Foundations of Intelligence. For the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: Peter Grindrod is an employee of Astut Ltd. The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Poston, T.; Stewart, I. *Catastrophe Theory and Its Applications*; Pitman: Boston, MA, USA, 1978.
2. Rosser, J.B. The rise and fall of catastrophe theory applications in economics: Was the baby thrown out with the bathwater? *J. Econ. Dyn. Control* **2007**, *31*, 3255–3280. [CrossRef]
3. Barunik, J.; Vosvrda, M. Can a stochastic cusp catastrophe model explain stock market crashes? *J. Econ. Dyn. Control* **2009**, *33*, 1824–1836. [CrossRef]
4. Klose, A.K.; Karle, V.; Winkelmann, R.; Donges, J.F. Emergence of cascading dynamics in interacting tipping elements of ecology and climate. *R. Soc. Open Sci.* **2020**, *7*, 200599. [CrossRef] [PubMed]
5. Gladwell, M. *The Tipping Point: How Little Things Can Make a Big Difference*; Little, Brown and Company: New York, NY, USA, 2000.
6. Arthur, W.B. Competing Technologies, Increasing Returns, and Lock-In by Historical Events. *Econ. J.* **1989**, *99*, 116–131. [CrossRef]
7. Arthur, W.B. Self-Reinforcing Mechanisms in Economics. In *The Economy as an Evolving Complex System*; Arthur, W.B., Durlauf, S.N., Lane, D.A., Eds.; Addison-Wesley: Reading, MA, USA, 1988; pp. 9–31.
8. Arthur, W.B. *Increasing Returns and Path Dependence in the Economy*; University of Michigan Press: Ann Arbor, MI, USA, 1994.
9. Stewart, I. *Does God Play Dice? The Mathematics of Chaos*; Basil Blackwell: Oxford, UK, 1989.
10. Stewart, I. *Game, Set and Math: Enigmas and Conundrums*; Penguin: Harmondsworth, UK, 1990.
11. Otter, N.; Porter, M.A.; Tillmann, U.; Grindrod, P.; Harrington, H.A. A roadmap for the computation of persistent homology. *EPJ Data Sci.* **2017**, *6*, 17. [CrossRef] [PubMed]
12. Edelsbrunner, H.; Letscher, D.; Zomorodian, A. Topological persistence and simplification. *Discret. Comput. Geom.* **2002** *28*, 511–533. [CrossRef]
13. Zomorodian, A.; Carlsson, G. Computing persistent homology. *Discret. Comput. Geom.* **2005**, *33*, 249–274. [CrossRef]
14. Carlsson, G. Topology and data. *Bull. Am. Math. Soc.* **2009**, *46*, 255–308. [CrossRef]
15. Gidea, M.; Katz, Y. Topological data analysis of financial time series: Landscapes of crashes. *Phys. A Stat. Mech. Its Appl.* **2018**, *491*, 820–834. [CrossRef]

16. Khasawneh, F.A.; Munch, E. Chatter detection in turning using persistent homology. *Mech. Syst. Signal Process.* **2016**, *70–71*, 527–541. [CrossRef]
17. Perea, J.A.; Harer, J. Sliding windows and persistence: An application of topological methods to signal analysis. *Found. Comput. Math.* **2015**, *15*, 799–838. [CrossRef]
18. Akingbade, S.W.; Gidea, M.; Manzi, M.; Nateghi, V. Why topological data analysis detects financial bubbles? *Commun. Nonlinear Sci. Numer. Simul.* **2024**, *128*, 107665. [CrossRef]
19. Guritanu, E.; Barbierato, E.; Gatti, A. Topological Machine Learning for Financial Crisis Detection: Early Warning Signals from Persistent Homology. *Computers* **2025**, *14*, 408. [CrossRef]

https://doi.org/10.3390/complexities2020010