26 RISKS
50,000 SIMULATED FUTURES

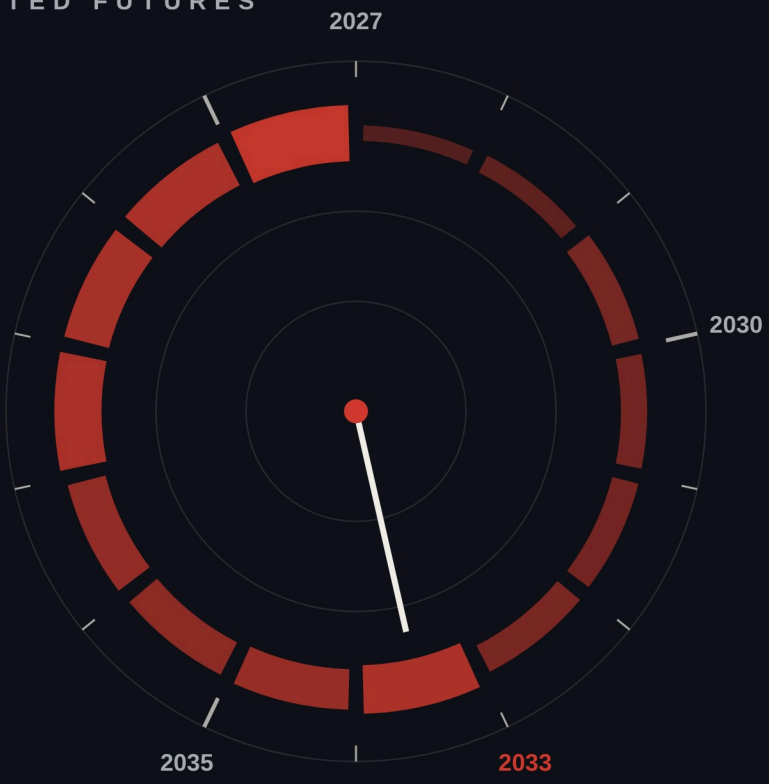

THE 2040 CLOCK

Probabilities, Timelines and Risk Maps of the Strategic Threats Facing America

2027 to 2040

GLOBAL STRATEGY FILES

THE 2040 CLOCK

Probabilities, Timelines and Risk Maps of the Strategic Threats Facing America

GLOBAL STRATEGY FILES

globalstrategyfiles.com

Copyright © 2026 Global Strategy Files. All rights reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in reviews and certain other noncommercial uses permitted by copyright law.

This book is a collective work of the contributors to Global Strategy Files. The probabilities it presents are model estimates produced with the method described in Chapter 2 and Annex A. They are analytical judgments, not certainties, and they are not investment, legal or security advice.

First edition, November 2026. Data current as of November 15, 2026.

"The dogmas of the quiet past are inadequate to the stormy present."

Abraham Lincoln, Annual Message to Congress, December 1, 1862

Contents

Section	Title
Introduction	The Clock Is Already Running
Part I	**The Method**
Chapter 1	Why Strategic Forecasts Fail
Chapter 2	The Risk Engine: How the Probabilities in This Book Are Made
Chapter 3	Twenty Years of Signals: The 2006 to 2026 Trend Atlas
Part II	**The Material Base**
Chapter 4	The Food System: From Breadbasket to Import Dependence
Chapter 5	Water: The Colorado, the Ogallala and the End of Abundance
Chapter 6	Power: The Grid in the Age of Artificial Intelligence
Chapter 7	Critical Minerals: The Chokepoint Beijing Controls
Part III	**The Technology Frontier**
Chapter 8	The Silicon Shield and Its Cracks: Chips and Taiwan
Chapter 9	Artificial Intelligence: Race, Accident and Misuse
Chapter 10	Invisible Fronts: Cyber, Quantum and Space
Part IV	**The Adversaries**
Chapter 11	China: The Decade of Maximum Danger
Chapter 12	Russia: The Wounded Bear and the NATO Frontier
Chapter 13	Iran, North Korea and the Nuclear Threshold
Chapter 14	The Axis Effect: When Adversaries Coordinate
Part V	**The Home Front**
Chapter 15	Debt and the Dollar: The Fiscal Fuse

Section	Title
Chapter 16	Demography: Fewer Births, Older Voters, Scarcer Workers
Chapter 17	The Fraying Republic: Political Violence and Contested Power
Chapter 18	Climate Shocks: Disasters, Heat and the Insurance Retreat
Chapter 19	Pathogens: Pandemics and Biological Attack
Part VI	**Synthesis**
Chapter 20	Cascades: How Crises Chain Together
Chapter 21	The Risk Maps: Where, When and How Hard
Chapter 22	2040: Four Futures and a Resilience Agenda
Conclusion	Reading the Clock
Annexes	
Annex A	Model Specification and Parameters
Annex B	The Risk Register: Definitions and Resolution Criteria
Annex C	Data Sources for the Trend Atlas
Annex D	The State Strategic Exposure Index
Annex E	Glossary

List of Figures and Tables

- Figure 2.1. The structure of the GSF Strategic Risk Clock
- Table 2.1. The twenty-six risks and their resolution thresholds
- Figure 3.1. The trend atlas: fourteen strategic indicators, 2006 to 2026
- Table 3.1. Twenty years of change in numbers
- Figure 4.1. The US agricultural trade balance, 2006 to 2026
- Figure 5.1. Lake Mead elevation and its thresholds, 2006 to 2026
- Figure 6.1. The end of flat demand: US electricity demand, 2006 to 2025
- Figure 11.1. China's nuclear build-up: estimated warheads, 2006 to 2024
- Figure 15.1. Federal debt held by the public, 2006 to 2026
- Figure 15.2. Net interest on the federal debt, 2006 to 2026
- Figure 15.3. The dollar's share of world reserves, 2006 to 2026
- Figure 16.1. The US total fertility rate, 2006 to 2025
- Figure 17.1. Threat cases against members of Congress, 2017 to 2025

- Figure 18.1. Billion-dollar weather disasters in the United States, 2006 to 2025
- Figure 20.1. The cascade network
- Figure 20.2. Conditional risk: how one crisis raises the odds of another
- Figure 21.1. The risk matrix: probability by 2040 against impact
- Figure 21.2. The arrival clock: when each risk is most likely to strike
- Figure 21.3. The hazard heat map, 2027 to 2040
- Figure 21.4. Probability of at least one event, by domain
- Figure 21.5. The State Strategic Exposure Index
- Figure 21.6. The global map of chokepoints and flashpoints
- Figure 21.7. Compound crises: how many risks, and the worst single year
- Figure 22.1. Sensitivity: which assumptions move the total
- Table 22.1. Four futures for 2040
- Figure 22.2. The resilience agenda: probabilities with and without action

Abbreviations

Abbreviation	Meaning
AI	Artificial intelligence
CBO	Congressional Budget Office

Abbreviation	Meaning
COFER	Currency Composition of Official Foreign Exchange Reserves (IMF)
CPI	Consumer price index
DoD	Department of Defense
EPRI	Electric Power Research Institute
ERCOT	Electric Reliability Council of Texas
FAS	Federation of American Scientists
IEA	International Energy Agency
IMF	International Monetary Fund
LBNL	Lawrence Berkeley National Laboratory
NATO	North Atlantic Treaty Organization
NERC	North American Electric Reliability Corporation
NOAA	National Oceanic and Atmospheric Administration
OASI	Old-Age and Survivors Insurance trust fund
OECD	Organisation for Economic Co-operation and Development
PLA	People's Liberation Army
PPP	Purchasing power parity
TWh	Terawatt-hour
USDA	United States Department of Agriculture

Introduction: The Clock Is Already Running

In the first ten months of 2026, the United States went to war with Iran, a war that has kept the Strait of Hormuz largely closed to commercial shipping for most of the year, pushed oil prices toward a hundred dollars a barrel and lifted the price of nitrogen fertilizer by more than a third in a single month. Lake Mead fell to the lowest level in its history, and federal hydrologists warned that Lake Powell could drop below the elevation at which Glen Canyon Dam can still generate electricity. The Congressional Budget Office reported that federal debt held by the public had passed the size of the economy and would exceed its Second World War record by 2030. Threat cases against members of Congress reached nearly fifteen thousand, more than three times the level of 2017. And in November, voters handed both chambers of Congress to the opposition in one of the largest midterm swings in half a century.

None of these events was a surprise to people who watch the underlying trends. Each had been building for years, sometimes for decades. What surprised almost everyone

was their simultaneity: a war in the Gulf, a drought in the West, a fiscal warning in Washington and a wave of political anger, all inside the same year, each making the others harder to manage. The year 2026 was a preview of the kind of strategic environment that this book argues will define the United States until 2040: not one great threat, but many interacting ones, arriving on overlapping schedules.

This book is an attempt to put numbers on that environment. It asks three questions about each of the major strategic risks facing the United States between 2027 and 2040. How likely is it? When is it most likely to arrive? And what happens to the other risks when it does? The answers come from a model built for this book, the GSF Strategic Risk Clock, which combines twenty years of data with historical base rates and simulates fifty thousand possible paths for the country, year by year, until 2040. The result is not a prophecy. It is a map of probabilities, a timetable of danger and an inventory of the connections that turn a bad year into a dangerous decade.

Why another book about threats?

Washington does not lack threat assessments. The intelligence community publishes one every year. The Pentagon publishes reports on China's military power, the National Intelligence Council publishes long-range global trends, and think tanks produce a constant stream of scenarios and war games. Insurance companies, central banks and

the World Economic Forum publish risk reports of their own. Why add another?

Because almost all of these documents share three weaknesses, and together those weaknesses blind decision makers to the shape of the next fifteen years.

First, they rarely state probabilities. Official assessments speak of threats that are "likely", "increasing" or "of concern". These words are deliberately elastic. Studies of intelligence analysts have long shown that readers interpret the same word, such as "probable", to mean anything from a coin toss to near certainty. A threat described as serious can be a one-in-a-hundred chance or a nine-in-ten chance, and the policy response appropriate to each is entirely different. This book gives every risk a number, with an uncertainty range, and explains where the number comes from.

Second, they rarely state timing. A threat that is likely to materialize by 2040 but unlikely before 2035 calls for a different response from one that could strike next year. Most assessments speak of "the coming decade" or "the long term" without distinguishing between risks that are imminent and those that are maturing. This book estimates, for each risk, the year in which it is most likely to arrive and the window within which it will probably arrive if it arrives at all.

Third, they treat threats one at a time. Assessments are organized by region or by topic: China, Russia, cyber, climate, health. Each chapter is written by a different office

with its own expertise and its own vocabulary. But the most dangerous moments in history were moments when several shocks arrived together and fed each other. A blockade of Taiwan would not only be a military crisis. It would be a semiconductor crisis, a critical-minerals crisis, a cyber crisis, a market crisis and possibly a food-price crisis, all within weeks. The model in this book treats these connections explicitly and measures how much each crisis raises the odds of the others.

What the clock shows

The headline results can be stated in a few sentences, and the rest of the book explains and qualifies them.

The United States is very likely to face several serious strategic shocks between now and 2040. In the model's central estimate, about thirteen of the twenty-six risks it tracks will occur at least once. In fewer than one simulated future in twenty does the country face eight or fewer of them. In the worst single year of a typical path, four or five major events arrive in the same twelve months, and in about half of all paths at least one year sees five or more. The question for American strategy is therefore not whether the country will be tested, but whether it will be tested by several crises at once, and whether its institutions can absorb that.

The most likely risks are not the most dramatic ones. They are the slow-moving failures of systems that Americans take for granted. A failure of the Colorado River system, defined as Lake Powell falling below its

minimum power pool or Lake Mead falling below 1,000 feet, has a probability of about 60 percent by 2030 and above 90 percent by 2040 in the model. A destructive cyberattack that disrupts critical infrastructure for at least a million people for a day or more is estimated at 86 percent by 2040. A food-price shock in which grocery prices rise by ten percent or more in a year is estimated at 80 percent. A critical-minerals embargo severe enough to halt American production lines for three months or more is estimated at 78 percent. A blackout affecting ten million customers or more for at least a day is estimated at 73 percent.

The gravest risks are less likely but far from negligible. The model estimates a 41 percent probability that China imposes a blockade or quarantine on Taiwan lasting at least thirty days, or launches an assault, by 2040, with the window of maximum danger in the early 2030s. It estimates a 17 percent probability of direct combat between American and Chinese forces causing at least a hundred American deaths, and a 14 percent probability that a nuclear weapon is detonated in conflict somewhere in the world. These are not forecasts that war will happen. They are measures of how much risk the country is carrying, and they are high enough to justify serious investment in deterrence.

The home front carries risks as large as the foreign ones. A crisis in the Treasury market, defined as a ten-year yield above 6.5 percent or an emergency intervention to restore market functioning, is estimated at 69 percent by

2040. The probability that the dollar's share of world reserves falls below 50 percent is 59 percent. The probability of a contested transfer of presidential power in at least one of the next four elections is 47 percent. Major political violence, including the assassination of a national figure, is estimated at more than 90 percent over the period, a reflection of how frequent such events have already become.

What is new in this book

Readers who know the literature on strategic risk will find five features that, to our knowledge, have not been combined before.

A single model for twenty-six risks across eight domains. Food and water, energy, technology, adversaries, fiscal stability, social cohesion, climate and health are usually studied by different communities with different methods. Here they share one framework, one time horizon and one unit of measurement, which makes them comparable for the first time.

Resolution criteria for every forecast. Each risk is defined by a threshold that can be checked against public data: an elevation of a reservoir, a yield on a bond, a number of customers without power, a count of deaths. Annex B lists them all. A reader in 2030 or 2040 will be able to score this book, risk by risk, and see where it was right

and where it was wrong. Forecasting that cannot be scored cannot improve.

Time, not just probability. The model produces, for each risk, a distribution of arrival years. It can say not only that a risk has a 40 percent chance of occurring by 2040, but that it is front-loaded or back-loaded, that it peaks in 2031 or rises steadily toward 2040. Figure 21.2, the arrival clock, is the visual summary of this dimension.

Cascades measured, not asserted. Many analysts say that crises are interconnected. Few measure by how much. The model contains thirty-one explicit delayed links between risks, twenty-one immediate ones, and two layers of hidden common drivers that make bad years cluster together. Chapter 20 shows which crises are the great multipliers and which are mostly consequences.

A geography of exposure. Strategic risk is not evenly spread across the country. Chapter 21 presents the GSF State Strategic Exposure Index, which combines water stress, climate hazard, grid stress, the concentration of military and strategic targets, and dependence on agriculture. Texas, California, Arizona and Nevada emerge as the most exposed states, for different reasons that matter to how they should prepare.

What this book is not

Three clarifications will prevent misunderstanding.

This book is not a prediction that the worst will happen. A probability of 17 percent for a US-China war means that,

in more than four futures out of five, there is no such war. The purpose of estimating such probabilities is to act so that the high-risk futures become less likely, not to resign oneself to them.

This book is not partisan. Several of the risks it measures, such as fiscal stress, political violence and contested elections, are politically charged. The model treats them as it treats reservoirs and pathogens: as systems with measurable trends and historical base rates. Where judgment is required, the book says so and shows how the result would change under different assumptions.

This book is not a final word. The model is a tool, not an oracle. Its parameters are published in Annex A so that others can test, criticize and improve them. Its estimates will be wrong in some cases, and the resolution criteria are there precisely so that those errors can be found.

How the book is organized

Part I, The Method, explains why strategic forecasts so often fail (Chapter 1), how the GSF Strategic Risk Clock works (Chapter 2) and what twenty years of data from 2006 to 2026 reveal about the direction of change (Chapter 3).

Part II, The Material Base, examines the physical foundations of American power: the food system and its new dependence on imports (Chapter 4), water in the West (Chapter 5), an electricity grid facing the largest demand

surge in a generation (Chapter 6) and the critical minerals that China controls (Chapter 7).

Part III, The Technology Frontier, turns to the technologies on which military and economic power now depend: advanced semiconductors and Taiwan (Chapter 8), artificial intelligence (Chapter 9), and the invisible fronts of cyber, quantum computing and space (Chapter 10).

Part IV, The Adversaries, assesses the states that could challenge the United States by force: China (Chapter 11), Russia (Chapter 12), Iran and North Korea (Chapter 13), and the growing coordination among them (Chapter 14).

Part V, The Home Front, examines risks generated inside the United States: the fiscal trajectory and the dollar (Chapter 15), demography (Chapter 16), political cohesion (Chapter 17), climate shocks (Chapter 18), and pandemics and biological attack (Chapter 19).

Part VI, Synthesis, brings the pieces together: how crises cascade (Chapter 20), where and when risk concentrates (Chapter 21), and four possible futures for 2040, with a resilience agenda that the model shows could reduce the country's expected strategic losses (Chapter 22).

The **Conclusion** explains how readers can use the clock, update it and hold its authors to account. The **Annexes** give the full model specification, the definitions of all twenty-six risks, the data sources, the state index and a glossary.

The clock in the title is not a doomsday clock. It does not move toward midnight. It measures time: the time the country has before slow trends become sudden crises, and the time it still has to act.

PART I: THE METHOD

Before a single probability can be trusted, the reader needs to know how it was made and why earlier attempts to see the future so often went wrong. Part I explains the failures of strategic forecasting, describes the model built for this book and presents twenty years of data on the forces that will shape the next fifteen.

Chapter 1. Why Strategic Forecasts Fail

In the spring of 2006, an analyst asked to describe the strategic risks facing the United States over the next twenty years would have produced a list that looks reasonable even today: terrorism, nuclear proliferation by Iran and North Korea, the rise of China, energy dependence on the Middle East, pandemic influenza, the fiscal costs of an aging population. Most of these items were real. Yet the same analyst would almost certainly have missed the events that actually defined the period: a financial crisis that nearly destroyed the global banking system two years later; a shale revolution that turned the United States into the world's largest producer of oil and gas; a smartphone and social-media revolution that transformed politics as much as commerce; a pandemic caused not by influenza but by a coronavirus; a full-scale Russian invasion of Ukraine; and an attack on the United States Capitol by Americans trying to stop the certification of an election.

The point is not that the analyst was foolish. It is that the methods used to think about strategic risk systematically produce certain kinds of errors. This chapter identifies

five of them, shows how each appeared in the forecasts of the past twenty years and explains how the model in this book is designed to reduce them. It cannot eliminate them. No method can. But a method that knows its own weaknesses is better than one that does not.

1.1 The test of 2006

One way to measure the quality of strategic foresight is to compare what serious observers expected in 2006 with what happened by 2026. Table 1.1 does this for ten questions that a well-informed analyst might have been asked. The column on the left summarizes the prevailing expectation among official and expert sources at the time; the column on the right records the outcome.

Table 1.1. The test of 2006: expectation and outcome

Question in 2006	Prevailing expectation	Outcome by 2026
Will US oil and gas production keep declining?	Yes; import dependence would rise	No; the United States became the largest producer of both
Will China's economy surpass the US by 2026?	Possibly at market exchange rates	No at market rates; yes in purchasing-power terms and in R&D spending by 2024
Will a major US bank fail?	Very unlikely	The 2008 crisis brought down or forced the rescue of several of the largest
Will the next pandemic come from influenza?	Most likely H5N1 avian influenza	A coronavirus killed more than a million Americans
Will federal debt exceed the size of the economy?	Not before the 2040s	It reached 100 percent of GDP in 2020 and stayed there

Question in 2006	Prevailing expectation	Outcome by 2026
Will Russia fight a large conventional war in Europe?	Very unlikely	It invaded Ukraine in 2014 and again on a massive scale in 2022
Will a US election result be violently contested?	Almost inconceivable	The Capitol was attacked in January 2021
Will the US run an agricultural trade deficit?	No; agriculture was a reliable surplus	Record deficits from 2023 onward
Will Lake Mead fall below 1,050 feet?	A worst case for the late century	It fell below that level in 2022 and again in 2026
Will the United States fight Iran directly?	A possibility, mostly from the air	A war that began in February 2026 has kept the Strait of Hormuz largely closed for months

The pattern is striking. On almost every question, the consensus underestimated the size of change. The errors were not random: they ran in the direction of continuity. The world of 2026 differs from the expectations of 2006 not in having different problems, but in having the same problems arrive earlier, faster and in combinations nobody had considered.

1.2 Blind spot one: the tyranny of words

The first failure is linguistic. Strategic assessments communicate likelihood with words rather than numbers. The US intelligence community has, since the 2000s, tried to standardize those words, assigning ranges to terms such as "likely" and "very likely". But the standard is not followed outside the intelligence community, and even inside it readers remember the word and forget the range.

Words create two problems. They hide disagreement: two analysts can both call a war "possible" while one means five percent and the other forty. And they prevent learning: a forecast that a threat is "growing" can never be proven wrong, because almost any outcome is compatible with it. Organizations that cannot be wrong do not improve.

Research on forecasting accuracy has shown the alternative. Over the past fifteen years, large forecasting tournaments sponsored by the US intelligence research agency demonstrated that ordinary people trained in a few simple habits, and scored on their numerical predictions, could outperform professional analysts with access to classified information. The habits were not mysterious: break questions into parts, start from the frequency of similar events in the past, update gradually as evidence arrives and keep score. The most important of these habits is the last. Numbers are not more truthful than words, but they can be checked.

This book therefore gives every risk a probability and a resolution criterion. When it says that a mega-blackout has a 73 percent probability by 2040, it means a blackout affecting at least ten million customers for at least twenty-four hours, and it means that if one could run the next fourteen years many times, such a blackout would occur in about seventy-three runs out of a hundred.

1.3 Blind spot two: neglect of base rates

The second failure is psychological. When people estimate the probability of an event, they tend to start from the story of the case in front of them rather than from the frequency of similar cases. Psychologists call this base-rate neglect. Its consequences in strategic analysis are severe in both directions.

It makes dramatic but rare events seem more likely than they are. After every major terrorist attack, public estimates of the probability of another attack rise far above what the historical record supports. The same happens after a hurricane, a cyberattack or a mass shooting.

It makes frequent but undramatic events seem less likely than they are. Food-price shocks of ten percent or more in a year have occurred three times in the United States in the past fifty-five years. Corn yields have fallen fifteen percent or more below trend about once a decade. Blackouts affecting millions of customers have occurred repeatedly. Yet each of these events tends to be treated as an exception when it happens, rather than as the recurrence of a known pattern.

The remedy is to begin every estimate with what forecasters call the outside view: how often has this kind of event happened, over what period, among comparable cases? Only then does one adjust for the specific features of the present, the inside view. The model in this book follows this order. Every risk starts from a reference class and a historical frequency. The adjustments for current trends

come second, and they are made explicitly, so that the reader can see how much of each estimate is history and how much is judgment.

1.4 Blind spot three: the straight line

The third failure is the assumption that trends continue at their recent pace, or, its mirror image, that they will soon revert to normal. Both are forms of linear thinking, and both fail when systems approach thresholds.

The Colorado River illustrates the problem. For more than two decades, water managers treated the decline of Lake Mead as a gradual trend punctuated by recoveries in wet years. The year 2011 seemed to vindicate them: the lake rose more than forty-five feet in twelve months after an exceptional snowpack. But a reservoir is not a line. It has thresholds: the elevations that trigger shortage declarations, the minimum power pool below which turbines must stop, the dead pool below which water can no longer flow by gravity through the dam. As a system approaches such thresholds, the consequences of each additional foot of decline grow rapidly, and the room for error shrinks to nothing.

Federal debt has a similar structure. A rise from 35 percent of GDP in 2006 to 100 percent in 2026 was absorbed with little visible stress, because interest rates were falling for most of the period. But interest costs depend on both the stock of debt and the interest rate, and when rates rose after 2022, net interest jumped from about 1.5 percent of GDP to more than 3 percent in four years. The

relationship between debt and fiscal stress is not a line but a curve that steepens.

The model handles thresholds in two ways. For risks defined by a physical or financial threshold, such as reservoir elevations or bond yields, the hazard is linked to the distance from the threshold and to the trend that moves the system toward it. For risks driven by structural pressure, such as fiscal stress, the hazard is made to rise more than proportionally with the underlying driver. In both cases, the shape of the curve is an explicit assumption that the reader can examine.

1.5 Blind spot four: one threat at a time

The fourth failure is organizational. Governments, universities and think tanks are divided into offices that specialize in particular threats. Each office sees its own threat clearly and the connections to others dimly. The result is a stack of separate assessments that cannot be added up.

This matters because the most damaging crises of the past twenty years were compound crises. The pandemic of 2020 was a public-health crisis that became a fiscal crisis, a supply-chain crisis, a political crisis and, for some months, a crisis of the Treasury market itself, when the market for the world's safest asset briefly stopped functioning in March 2020 and required emergency purchases by the Federal Reserve. The war with Iran that began in February 2026 is a military crisis that has become an en-

ergy crisis, a fertilizer crisis and an inflation crisis, and it helped shape the outcome of the midterm elections.

Compound crises are not simply the sum of their parts. When several shocks arrive together, they compete for the same scarce resources: the attention of leaders, the capacity of emergency services, the fiscal space of the Treasury, the patience of the public. A country that could handle any one of them might be overwhelmed by three at once.

The model in this book is built around connections. It contains thirty-one delayed cascades and twenty-one immediate consequences, links through which one event raises the probability of another, either immediately or over the following years. It also contains three hidden common drivers, which represent conditions that raise many risks at the same time: the level of geopolitical tension, the degree of climate stress and the degree of political and fiscal strain. These drivers move from year to year, so that bad years tend to cluster, as they do in history.

1.6 Blind spot five: the comfort of the present

The fifth failure is cultural. Societies that have lived through long periods of stability come to see stability as the natural state of affairs and disruption as an anomaly. Analysts call this normalcy bias. It explains why the warnings that preceded many crises, from the 2008 financial collapse to the 2020 pandemic, were heard but not acted upon.

Normalcy bias is particularly strong in the United States, for understandable reasons. The country has not been invaded in living memory, its institutions have survived civil conflict and economic depression, and its economy has repeatedly recovered from shocks that devastated others. These are real strengths. But they also encourage the assumption that what has not happened will not happen: that the dollar will always be the world's reserve currency, that the Treasury market will always function, that elections will always produce an orderly transfer of power, that water will always flow from the tap.

The remedy is not pessimism. It is the discipline of asking, for each assumption of stability, what would have to happen for it to fail, and how often such conditions have occurred elsewhere or in the past. The model asks this question for each of the twenty-six risks. The answer is sometimes reassuring: a cryptographically relevant quantum computer, for example, faces engineering obstacles that make its arrival before the early 2030s unlikely. Sometimes it is not.

1.7 What a good strategic forecast looks like

From these five failures, it is possible to derive seven standards that a useful strategic forecast should meet. They are the design principles of the model presented in the next chapter.

1. **Quantified.** Every risk has a probability, not a word.
2. **Resolvable.** Every risk is defined by a threshold that can be checked against public data, so that the forecast

can be scored.

3. **Grounded in base rates.** Every estimate starts from the historical frequency of comparable events.
4. **Trend-aware.** Adjustments for current conditions are explicit and tied to measured trends.
5. **Time-resolved.** The forecast says not only how likely a risk is by 2040 but when it is most likely to arrive.
6. **Connected.** The forecast models the links between risks and the common drivers that make crises cluster.
7. **Honest about uncertainty.** The forecast reports ranges, not just central estimates, and shows how sensitive its results are to its assumptions.

No forecast meets these standards perfectly. Some of the risks in this book, such as a catastrophic incident involving artificial intelligence, have no historical base rate at all, and the estimates for them rest heavily on judgment. The book says so wherever it applies. But a forecast that tries to meet these standards can be criticized, tested and improved. A forecast that does not try cannot.

1.8 Conclusion

The history of strategic foresight is mostly a history of surprise. The surprises were rarely complete: someone, somewhere, usually warned of each crisis before it happened. The failure lay in how warnings were expressed, weighed and connected. They were expressed in words that could mean anything, weighed against stories rather than frequencies, projected along straight lines toward

thresholds, studied one at a time and dismissed by a culture that had come to expect stability.

The model presented in the next chapter is an attempt to correct these failures systematically. It will not foresee the future. But it can show, more clearly than words alone, how much risk the United States is carrying, when that risk is concentrated and how its pieces fit together.

Notes to Chapter 1

- On estimative language and its interpretation by readers: Sherman Kent, "Words of Estimative Probability" (Central Intelligence Agency, Studies in Intelligence, 1964); Office of the Director of National Intelligence, Intelligence Community Directive 203, Analytic Standards.
- On forecasting tournaments and the habits of accurate forecasters: Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (Crown, 2015).
- On base rates and the outside view: Daniel Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux, 2011).
- Outcome data in Table 1.1: US Energy Information Administration; OECD Main Science and Technology Indicators (March 2026); Congressional Budget Office historical budget data; US Department of Agriculture Economic Research Service; US Bureau of Reclamation.

Chapter 2. The Risk Engine: How the Probabilities in This Book Are Made

Every number in this book comes from a single model, the GSF Strategic Risk Clock. This chapter explains how it works in terms that do not require a background in statistics. Readers who want the equations and parameter values will find them in Annex A; readers who want the precise definition of each risk will find it in Annex B. The aim here is to make the logic transparent enough that a reader can judge whether to trust the results, and to see exactly where judgment enters.

The model can be described as an answer to a simple question asked fifty thousand times: what might happen to the United States, year by year, between 2027 and 2040? Each answer is a path, a possible history of the next fourteen years in which some risks occur, some do not, and those that occur change the odds of others. Counting across all fifty thousand paths gives the probabilities, timings and connections reported throughout the book.

2.1 Step one: defining the risks

A forecast is only as good as the question it answers. The model tracks twenty-six risks, chosen by three criteria. First, each must be capable of causing strategic damage to the United States, meaning damage that affects national security, economic stability or the functioning of government on a national scale, not merely local harm. Second, each must be defined by a threshold that can be observed in public data, so that a reader in 2040 can determine whether it happened. Third, the set as a whole must cover the eight domains that, in our judgment, contain the main sources of strategic risk: food and water, energy, technology, adversaries, fiscal stability, social cohesion, climate and health.

The thresholds are deliberately high. A drought is not a risk in this model; a drought that cuts the national corn yield fifteen percent below trend is. A cyberattack is not a risk; a destructive cyberattack that disrupts critical infrastructure for at least a million people for at least a day is. A rise in Treasury yields is not a risk; a ten-year yield above 6.5 percent, or an emergency intervention by the authorities to restore market functioning, is. Setting the bar high has two advantages. It focuses attention on events that would genuinely matter at the national level, and it makes resolution unambiguous: there may be argument about whether a given cyberattack was "serious", but there is little argument about whether a million people lost power for a day.

Table 2.1. The twenty-six risks of the GSF Strategic Risk Clock

Domain	Code	Risk and resolution threshold
Food and water	FOOD_PRICE	Food price shock: food-at-home consumer prices up 10 percent or more year on year
Food and water	BREADBASKET	Breadbasket failure: US corn yield 15 percent or more below trend
Food and water	COLORADO	Colorado River system failure: Lake Powell below 3,490 feet or Lake Mead below 1,000 feet
Food and water	AQUIFER	Ogallala irrigation collapse: 25 percent or more of irrigated acres in Texas or Kansas lost
Energy	GRID	Mega-blackout: 10 million or more customers without power for 24 hours or more
Energy	GEOMAG	Extreme geomagnetic storm causing physical damage to the grid
Technology	MINERALS	Critical-mineral embargo halting some US production lines for 3 months or more
Technology	CHIPS	Loss of 50 percent or more of Taiwan's advanced-chip output for 3 months or more
Technology	CYBER	Destructive cyberattack on critical infrastructure affecting 1 million people or more for 24 hours or more
Technology	AI_INCIDENT	Catastrophic AI incident: $100 billion or more in losses, or 100 or more deaths
Technology	QDAY	A cryptographically relevant quantum computer is demonstrated
Adversaries	TAIWAN	Chinese blockade or quarantine of Taiwan lasting 30 days or more, or an assault
Adversaries	US_CHINA_WAR	Direct combat between US and Chinese forces with 100 or more US deaths
Adversaries	RUSSIA_NATO	Direct armed clash between Russian and NATO forces with 10 or more deaths

Domain	Code	Risk and resolution threshold
Adversaries	NUCLEAR_USE	Nuclear detonation in conflict anywhere in the world
Adversaries	IRAN_BOMB	Iranian nuclear test or declared nuclear weapon
Adversaries	HOMELAND	Mass-casualty attack on US soil with 100 or more deaths
Fiscal	FISCAL	Treasury market crisis: emergency intervention or ten-year yield of 6.5 percent or more
Fiscal	DOLLAR	Dollar share of world foreign-exchange reserves falls below 50 percent
Fiscal	SOCSEC	Social Security trust-fund cliff: automatic benefit cut or emergency deal under deadline
Society	CERTIFICATION	Contested transfer of presidential power
Society	POLVIOLENCE	Major political violence: assassination of a national figure or an attack with 10 or more deaths
Climate	MEGA_DISASTER	A single weather disaster costing $200 billion or more
Climate	HEAT	Lethal heat wave: 2,000 or more US deaths in a single week
Health	PANDEMIC	Pandemic causing 100,000 or more US deaths in a single year
Health	BIO_ATTACK	Deliberate biological attack causing 1,000 or more casualties

Several risks that readers might expect are absent, and the omissions are deliberate. A recession is not included because it is too frequent to be strategic in itself; the model captures the strategic version, a crisis in the Treasury market. Terrorism in general is not included, only mass-casualty attacks on American soil. Migration is not modeled as a risk in itself, because its effects are bet-

ter captured through its political consequences. Some risks, such as an asteroid impact or a supervolcanic eruption, are excluded because their probabilities over fourteen years are so small that they would not affect any result in the book.

2.2 Step two: the base rate

For each risk defined by a constant or slowly changing hazard, the model begins with a base rate: the median annual probability that the risk occurs in 2027, the first year of the simulation. This number is called the baseline hazard. It is estimated, wherever possible, from the frequency of comparable events in the historical record.

Some examples show how this works. Pandemics causing at least 100,000 American deaths in a year have occurred four times since 1900: in 1918, 1957, 1968 and 2020. Four events in about 125 years gives an annual frequency of roughly 3.2 percent, and the model's baseline hazard of 3 percent sits just below it. Food-at-home price inflation of 10 percent or more has occurred in 1973, 1974, 1978 to 1980 and 2022, which, counting episodes rather than years, gives about three episodes in fifty-five years, or roughly 5 percent a year; the model uses 4.5 percent. Corn yields at least 15 percent below trend occurred in 1983, 1988, 1993 and 2012, four times in about seventy years, or 5.7 percent a year; the model uses 5.5 percent.

For risks with no historical precedent, such as a catastrophic accident involving artificial intelligence, no frequency can be counted. In these cases the model uses the

best available proxies, such as the frequency of large losses from new technologies or the judgments of expert surveys, and it gives them a much wider uncertainty range. The book flags these cases explicitly. They are the parts of the model where the reader should hold the numbers most lightly.

2.3 Step three: uncertainty about the base rate

A base rate is itself uncertain. Four pandemics in 125 years could reflect a true annual rate of 2 percent or 6 percent; the record is too short to say. The model represents this by treating each baseline hazard not as a single number but as a range. In technical terms, it draws the baseline hazard from a lognormal distribution around the median, with a spread that reflects how much the evidence constrains it.

This matters more than it might seem. When the uncertainty about a hazard is large, the probability that the event occurs at least once over fourteen years is not simply what one would compute from the median. Some simulated worlds will be high-hazard worlds, in which the event is very likely; others will be low-hazard worlds, in which it rarely occurs. The model captures this by drawing two thousand different parameter sets, each representing a plausible version of the world, and running twenty-five simulations in each. The resulting fifty thousand paths contain both the uncertainty about which world we live in and the randomness of events within any given world.

2.4 Step four: trends over time

The world of 2040 will not have the same hazards as the world of 2027. Some risks are rising because their drivers are growing; others are expected to peak and decline. The model represents this with a time profile for each risk, a multiplier applied to the baseline hazard that changes from year to year. Each profile is tied to a measured trend from the data presented in Chapter 3, and each has its own uncertainty.

There are nine profiles in the model.

- **Flat.** The hazard stays at its baseline level. Used where no clear trend exists, such as for geomagnetic storms, which depend on the sun's activity rather than on human behavior.
- **Slow rise.** The hazard grows by about 2.5 percent a year, with uncertainty of about one percentage point. Used for risks whose drivers are growing steadily but slowly, such as food-price shocks, pandemics and political violence.
- **Fast rise.** The hazard grows by about 6 percent a year. Used for risks whose drivers are accelerating, such as cyberattacks, lethal heat waves and the depletion of the Ogallala aquifer.
- **Capability growth.** The hazard grows by about 15 percent a year, capped at eight times its starting value. Used for risks linked to the rapid improvement of artificial intelligence, including catastrophic AI incidents and biological attacks enabled by new tools.

- **Colorado.** A high near-term hazard that declines partly as new operating guidelines for the river take effect after 2026, but with a slow upward drift reflecting warming and declining runoff.
- **Grid.** A hump that peaks around 2031, when the surge in electricity demand from data centers is expected to outpace new generation and transmission, followed by a partial easing as the build-out catches up.
- **Minerals.** A high hazard that begins to decline after 2030 as non-Chinese sources of refining come online, reducing Beijing's leverage.
- **Taiwan.** A hump whose peak falls somewhere between 2028 and 2033, the period that many analysts consider the window of maximum danger, with an amplitude that varies between 1.4 and 2.6 times the baseline.
- **Fiscal.** A hazard tied to the ratio of federal debt to GDP along the path projected by the Congressional Budget Office, rising more than proportionally as debt grows from 101 percent of GDP in 2026 toward roughly 139 percent in 2040.

The last of these deserves a word of explanation, since it illustrates how the model handles nonlinearity. The fiscal hazard is proportional to the debt ratio raised to a power between two and four. If the power were one, a 38 percent increase in the debt ratio would increase the hazard by 38 percent. With a power of three, it roughly multiplies it by 2.6. The choice reflects the evidence from other advanced economies that market tolerance of debt is not linear: it

holds for a long time and then gives way quickly, often triggered by an unrelated shock.

2.5 Step five: the hidden common drivers

If each risk were independent of all the others, the model could stop here. But risks are not independent. Wars, droughts, recessions and political crises tend to cluster in time, because they share underlying causes. The 1970s saw an oil embargo, a food-price shock, high inflation, political scandal and military defeat in the same decade. The years from 2020 to 2022 saw a pandemic, a crisis in the Treasury market, an attack on the Capitol, a war in Europe and the highest inflation in forty years.

The model represents these shared causes with three hidden factors, which statisticians call latent variables. Each moves from year to year, with memory: a high value this year makes a high value next year more likely. Each also has a slight upward drift, reflecting the judgment that the underlying pressures are growing.

- **Geopolitical tension (G)** captures the general level of hostility among great powers. When it is high, the risks of a Taiwan crisis, a Russia-NATO clash, a critical-minerals embargo and a destructive cyberattack all rise together.
- **Climate stress (C)** captures the severity of the weather in a given year. When it is high, the risks of a breadbasket failure, a mega-disaster, a lethal heat wave, a blackout and a Colorado River failure all rise together.

- **Political and fiscal strain (P)** captures the degree of polarization, institutional stress and fiscal pressure in American politics. When it is high, the risks of political violence, a contested election and a Treasury market crisis all rise together.

Each risk has a loading on each factor, a number between zero and one that measures how strongly it responds. A Taiwan crisis has a loading of 1.0 on geopolitical tension and 0.1 on political strain. A breadbasket failure has a loading of 0.8 on climate stress and nothing else. A food-price shock responds to all three: to geopolitics through trade disruption, to climate through harvests and to political strain through policy.

One technical detail matters for the results. Multiplying a hazard by a random factor that is sometimes high and sometimes low raises the average hazard, because the high values have a larger effect than the low ones. Without correction, adding common drivers would make every risk more likely, not just more clustered. The model applies a standard correction so that, on average, the factors leave each risk's probability unchanged and only alter how risks move together. This keeps the effects of clustering separate from the effects of the base rates.

2.6 Step six: cascades

The common drivers make risks cluster. Cascades make them chain. A cascade is a direct causal link through which the occurrence of one risk raises the probability of another. The model contains two kinds.

Delayed cascades raise the hazard of the target risk for a number of years after the trigger. A blockade of Taiwan multiplies the annual hazard of a destructive cyberattack by four for the following two years, of a critical-minerals embargo by four for three years and of a Treasury market crisis by three for two years. A Treasury market crisis raises the hazard of major political violence by 40 percent for two years. A geomagnetic storm multiplies the hazard of a mega-blackout thirty-fold for a year, reflecting the physical damage such a storm would do to transformers. There are thirty-one delayed cascades in the model.

Immediate consequences capture events that would follow the trigger within the same year, with a stated conditional probability. If China blockades Taiwan, there is a 55 percent chance that Taiwan's advanced-chip output is cut by half for three months or more in the same year, a 40 percent chance of a critical-minerals embargo, a 30 percent chance of a destructive cyberattack on American infrastructure and a 22 percent chance of direct US-China combat. If direct combat occurs, there is an 85 percent chance of a loss of Taiwan's chip output, a 65 percent chance of a destructive cyberattack and a 35 percent chance of a Treasury market crisis. There are twenty-one immediate consequences in the model.

The values come from three sources: historical analogues, such as the effect of the 1973 oil embargo on food prices; war games and scenario studies published by defense research institutions; and structural reasoning about the systems involved. They are among the most judgment-

laden parameters in the model, and Chapter 22 shows how sensitive the results are to them. But leaving cascades out would not make the model more objective. It would simply build in the assumption that the cascades are zero, which is certainly wrong.

2.7 Step seven: the special processes

Four risks do not fit the hazard framework and are modeled separately, each in the way that best fits its nature.

Q-Day, the arrival of a quantum computer capable of breaking the public-key cryptography that protects most of the world's communications, is not an event that happens by chance in a given year. It is a technological milestone that will arrive at some point. The model represents it as a time-to-event distribution whose median lies somewhere between 2035 and 2045, consistent with the range of expert surveys, with a shape that makes very early arrival unlikely.

The dollar's share of world reserves is a continuous quantity that moves gradually. The model starts it at 57.1 percent, the level reported by the International Monetary Fund for the first quarter of 2026, and lets it follow a random path with a downward drift of about half a percentage point a year, the average pace of the past two decades, plus random annual variation. The risk occurs when the share falls below 50 percent.

The Social Security trust-fund cliff is scheduled by law. Under current projections, the reserves of the Old-Age

and Survivors Insurance trust fund will be depleted in 2033 or 2034, at which point benefits would be cut automatically to the level that incoming revenue can support. The model assigns a 70 percent chance to depletion in 2033 and 30 percent to 2034, and a 35 percent chance that Congress enacts a fix before the deadline, between 2028 and 2032. If no early fix occurs, the risk is realized in the year of depletion, either through an automatic cut or through an emergency deal reached under the pressure of the deadline.

A contested transfer of presidential power can only occur in a presidential transition year. The model therefore evaluates it only in 2029, 2033, 2037 and, for the final certification of the 2040 election, in 2040 itself, though strictly that certification takes place in January 2041. In each of these years the probability is drawn from a range centered on 13 percent and is raised or lowered by the level of political strain in that year. The base value reflects the judgment that two of the last seven transitions, in 2001 and 2021, were seriously contested, though only one involved an attempt to prevent certification by force.

2.8 Step eight: simulation and counting

With all the pieces in place, the model runs. For each of fifty thousand paths, it moves year by year from 2027 to 2040. In each year it computes the hazard of every risk, combining the baseline, the time profile, the common drivers and any cascades triggered by earlier events. It then draws, for each risk, whether it occurs that year.

Immediate consequences are applied next, and the cascades they trigger are recorded for future years. The special processes are updated. The model then moves to the next year.

When all paths are complete, counting gives the results:

- the **probability by year**, the share of paths in which a risk has occurred at least once by a given year, reported in this book for 2030, 2035 and 2040;
- the **arrival distribution**, the years in which the risk first occurs in those paths where it occurs at all, summarized by its median and by the tenth and ninetieth percentiles;
- the **conditional probabilities**, the share of paths with a given risk that also contain another, compared with the share of all paths that contain it, which gives the lift reported in Chapter 20;
- the **compound statistics**, such as the number of distinct risks realized in each path and the number of events in the worst year.

2.9 Measuring impact

Probability is half of risk. The other half is consequence. The model scores each risk on three dimensions of impact, each on a scale of one to ten:

- **Human impact**: deaths, injuries and displacement;
- **Economic impact**: losses to output, wealth and public finances;

- **Strategic impact**: damage to American military power, alliances, global standing and the functioning of government.

The composite impact is a weighted average: 40 percent human, 35 percent economic and 25 percent strategic. The weights reflect a judgment that human lives come first, but that a book about strategic risk must also give weight to the foundations of national power. Different weights would change the ranking at the margins, and Annex A reports the main alternatives. A nuclear detonation scores ten on every dimension and has a composite of ten. A critical-minerals embargo scores one for human impact but seven for strategic impact, giving a composite of about 4.3.

Multiplying the probability by 2040 by the composite impact gives the **expected loss** for each risk, a single number that allows risks of very different kinds to be compared. Expected loss is not a prediction of damage; it is a way of ranking where attention and resources would do the most good. A risk with a 90 percent probability and moderate impact can rank alongside one with a 20 percent probability and severe impact. Chapter 21 presents the full ranking.

2.10 What the model cannot do

Every model simplifies, and it is important to be clear about the main simplifications in this one.

It counts first occurrences. Most results describe whether a risk occurs at least once by a given year, not how many times. A second Taiwan crisis or a third mega-blackout would matter, but the first occurrence is the one that tests the country's preparation.

It does not model policy responses explicitly, except in the resilience scenario of Chapter 22. The baseline assumes that current policies continue and that the trends of the past twenty years persist with the uncertainty shown. In reality, governments will respond to warning signs, and some of those responses will reduce risk. The model's central estimates are best read as the risk the country carries if it does not change course.

It does not include unknown unknowns. By construction, the model can only track risks that have been defined in advance. History suggests that some of the most important events of the next fourteen years will not be on anyone's list. The true probability of a severe strategic shock is therefore higher than the model's figures alone suggest.

It rests on judgment for the parameters that history cannot supply. The base rates are anchored in data wherever possible, but the cascades, the factor loadings and the hazards for unprecedented risks rely on structured judgment. These parameters are published in Annex A so that others can challenge them.

Finally, it is a model of the United States as a whole. It does not distinguish between a blackout in Texas and a

blackout in New England, except in the state-level analysis of Chapter 21, which uses a separate index rather than the simulation.

2.11 How to read the numbers

Throughout the book, probabilities are reported to the nearest whole percent. This should not be mistaken for precision. The uncertainty around most estimates is several percentage points, and for the least-anchored risks it is much more. A useful rule of thumb is to treat differences of less than five percentage points between risks as insignificant, and to focus instead on orders of magnitude and on the shape of the timing.

The distinction between probability and timing is the heart of the model. A risk with a probability of 40 percent by 2040 and a median arrival year of 2031, such as a Taiwan crisis, demands attention now. A risk with a similar probability but a median arrival of 2036, such as a catastrophic AI incident, demands preparation now and action that intensifies over the decade. A risk with a probability above 90 percent and a median arrival of 2029, such as a Colorado River failure, is no longer a risk to prevent but an outcome to manage.

Figure 2.1 summarizes the architecture of the model. The next chapter presents the data on which its time profiles are built.

Figure 2.1. The structure of the GSF Strategic Risk Clock

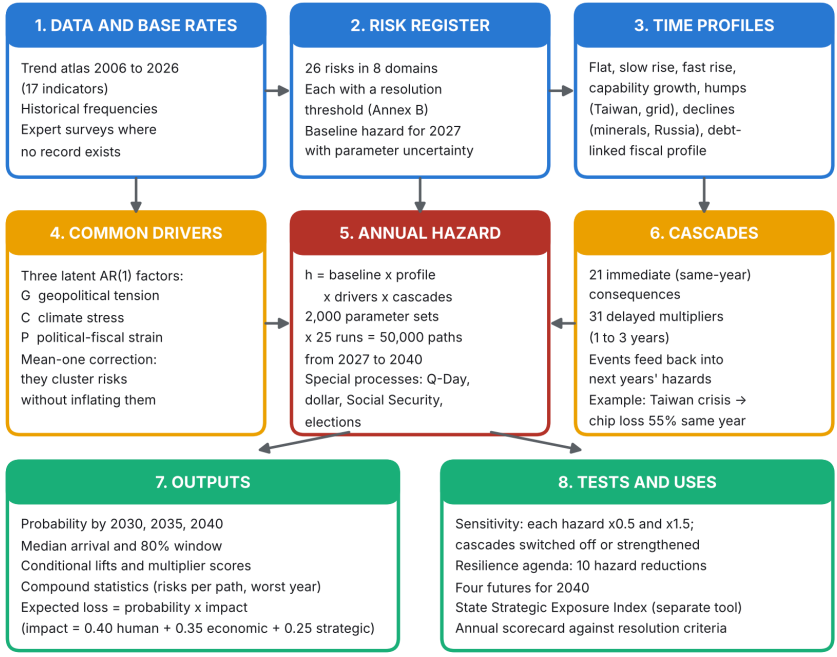

Source: GSF Strategic Risk Clock, model architecture (Annex A).

Notes to Chapter 2

- Pandemic mortality: Centers for Disease Control and Prevention, historical influenza pandemic estimates; National Center for Health Statistics, COVID-19 mortality data.
- Food prices: US Bureau of Labor Statistics, consumer price index for food at home, 1970 to 2025.
- Corn yields: US Department of Agriculture, National Agricultural Statistics Service, yield series 1955 to 2025, detrended by the authors.
- Fiscal trajectory: Congressional Budget Office, *The Budget and Economic Outlook: 2026 to 2036* and *The Long-*

Term Budget Outlook (2026).

- Dollar reserve share: International Monetary Fund, Currency Composition of Official Foreign Exchange Reserves, first quarter of 2026.
- Social Security: Social Security Board of Trustees, 2026 Annual Report.
- Quantum computing timelines: Global Risk Institute, *Quantum Threat Timeline Report* (successive editions).

End of the free excerpt

The complete book, with twenty-two chapters in six parts, a conclusion, ten annexes, thirteen original figures and risk maps, the full parameters of the GSF Strategic Risk Clock and the State Strategic Exposure Index for every state, is available as a PDF at globalstrategyfiles.com.

THE 2040 CLOCK

What are the odds that China blockades Taiwan before 2040? That the Colorado River system fails, the Treasury market seizes up, a cyberattack darkens a region or a presidential transition is contested? And what happens to everything else when one of them occurs?

Built on twenty years of data, the GSF Strategic Risk Clock simulates 50,000 possible paths for the United States from 2027 to 2040. It estimates the probability of twenty-six strategic risks, the year each is most likely to arrive and the cascades that turn one crisis into several. About thirteen risks occur in a typical future; in more than half of all futures, one year brings five or more major shocks at once, with 2033 the most dangerous year of all. The largest expected losses come from the foundations of national life: the cyber domain, the grid, Western water, the Treasury market and the peaceful transfer of power. A Taiwan crisis, estimated at 41 percent by 2040, is the hinge on which everything else turns.

With original risk maps, an arrival clock, a cascade network, a state-by-state exposure index, four scenarios for 2040 and a ten-point resilience agenda.

ABOUT THE AUTHORS

Global Strategy Files is an independent analysis platform devoted to geopolitics, strategic foresight and the forces reshaping international security and the world economy. Its contributors combine expertise in international relations, defense, economics, energy, technology and quantitative risk analysis. *The 2040 Clock* is a collective work of the site's contributors.