IABG

THREE TRANSFORMATIVE ENHANCEMENTS FOR LCC MODELS

Estimating for Life

Fabian Eilingsfeld, Birgit Biehn, Christian Schindler
ICEAA International Training Symposium
London, 19–20 September 2024

Lecture Outline

0 Life Cycle Cost Management (LCCM) in the Bundeswehr

1 “Bathtub Curve”: The cost of increased failure rates

2 “No Fault Found”: The cost of condition-based maintenance

3 “Parts Obsolete”: The cost of product and technology obsolescence

4 Lessons Learned and Conclusion

0

Estimating for Life.

Life Cycle Cost Management (LCCM)
in the Bundeswehr

IABG SERVICE PORTFOLIO

Life Cycle Cost Management (LCCM)

Cost forecasts and economic feasibility studies for public sector and industry:

- Life cycle cost management according to Bundeswehr directive A-1510/1
- Performance studies according to federal accounting directive
- Analyses (financial needs, sensitivity, scenarios, cost and schedule risk)

Various methods and tools:

- Parametric cost estimation
- Profitability analysis
- Analytical cost calculation
- Cost logging and data recording during operation and support (O&S)
- Monte Carlo simulations
- Spare parts cost simulations

The Bundeswehr began Life Cycle Cost Management (LCCM) in 2012 to optimise the economic performance of its programme portfolio

In systems procurement, LCCM has become the law

"According to the Life Cycle Cost Management Directive (A-1510/1), LCCM is a business management method that involves the use of life cycle costs (LCC) for decisions in the cycle of products or services." (Bundeswehr, 2015)

LCCM serves to "reinforce the economic focus of decision-making within the Federal Ministry of Defence." (Bundeswehr, 2022)

Bottom Line

The aim of LCCM is to optimise the life cycle costs (LCC) of a system, from development and procurement through to operation and support (O&S).

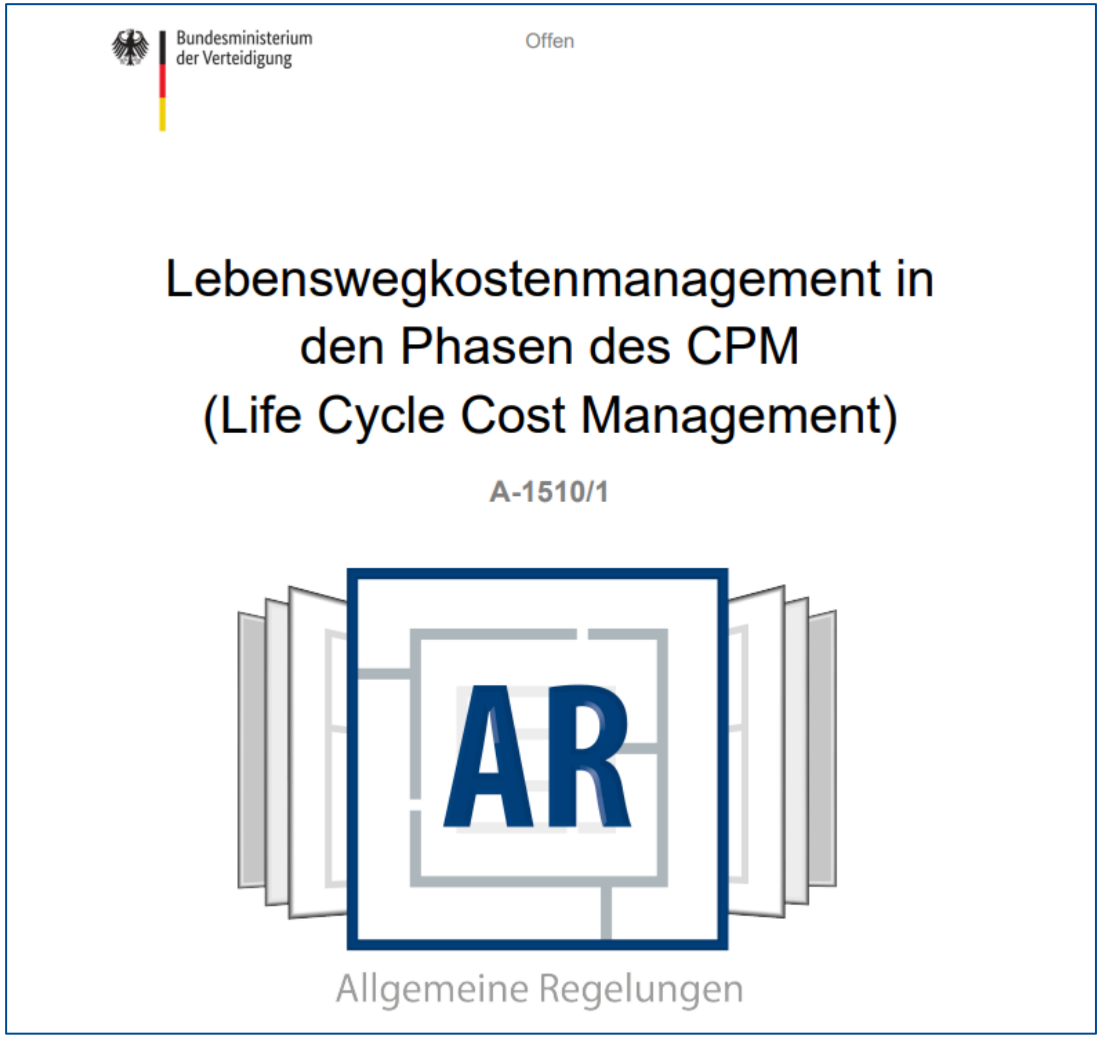

Bundesministerium
der Verteidigung

Offen

Lebenswegkostenmanagement in
den Phasen des CPM
(Life Cycle Cost Management)

A-1510/1

Allgemeine Regelungen

With LCCM, the Bundeswehr provides a rationale for enhancing existing life cycle cost models

Situation: Life Cycle Cost (LCC) is the basis for programme decisions

Life Cycle Cost Management (LCCM) starts with the design and planning of new systems. Comprehensive cost estimates and analyses shall be conducted early to ensure that long-term financial implications are considered during decision-making.

Before selecting a winner, life cycle costs for each proposed solution are planned.

"Planned" cost involves appropriation of line items for the federal budget, so cost uncertainty shall be as low as possible!

Challenge: O&S cost is the biggest part of LCC, but relatively uncertain

Operation and support (O&S) costs are inherently uncertain and challenging to estimate due to the influence of numerous variable factors that are difficult to predict.

These costs are subject to fluctuations over time, driven by market dynamics, inflation, technological advancements, political decisions, and other external factors.

Better models are needed for estimating O&S cost.

Solution: Quantify and model the impact of "known unknowns" on LCC

Several factors affecting the cost of a system's operation and support (O&S) are recognized as *known unknowns*. This term indicates that while we can anticipate their occurrence and impact on life cycle costs, their exact timing and magnitude remain uncertain.

Despite longstanding recognition of these influential cost drivers, traditional cost models have often overlooked them.

Life Cycle Cost models will improve if they take into account existing known unknowns.

In traditional Life Cycle Cost Models, the Earth is flat

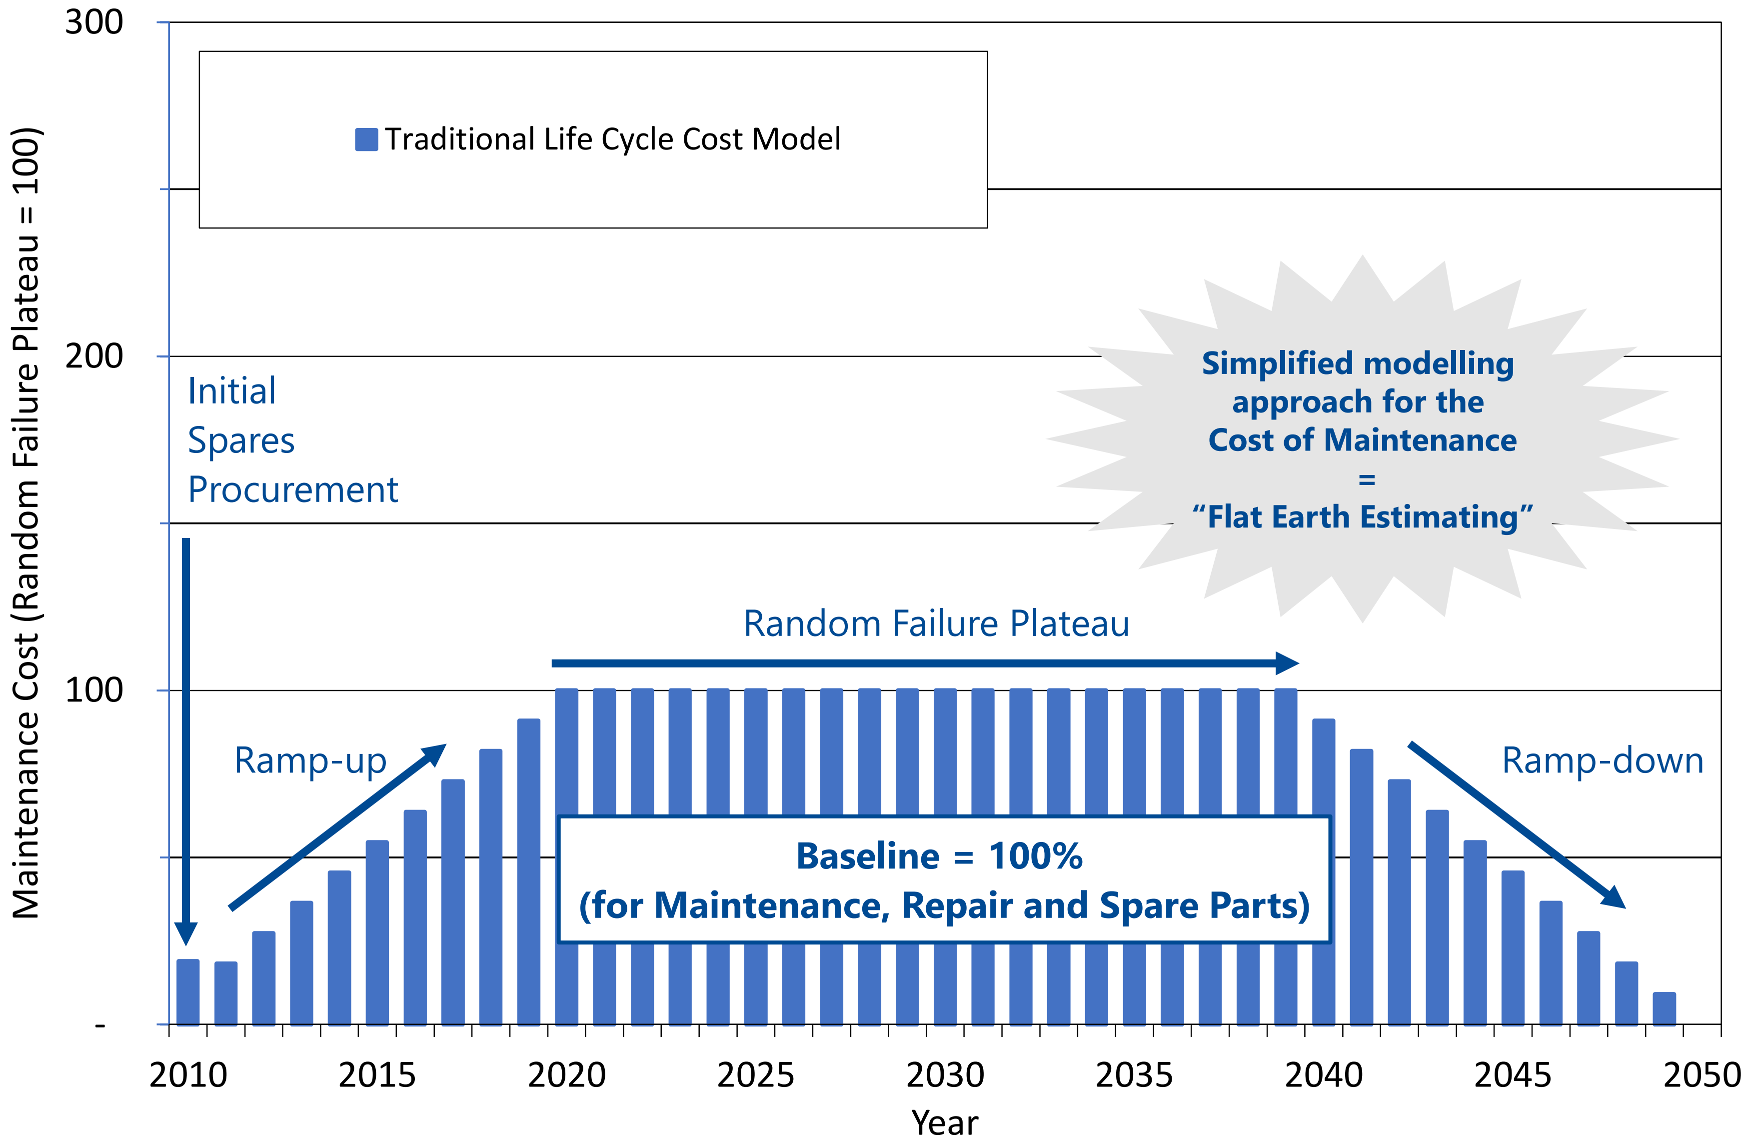

Life cycle costing was established in the 1970s. Philosophy in those days was to base operation and support (O&S) cost on constant failure rates, following the exponential failure law.
With maintenance intervals constant, established estimating practice turned out **constant maintenance cost** over time (without inflation).

Units have constant failure rates; all failures are random.

Units only undergo maintenance after suffering functional failure (corrective maintenance).

Spares never change over the entire service life and stay technically identical to the units from the initial procurement.

1

Non-Constant Failure Rates

“Bathtub Curve”:
The cost of increased failure rates

Introduction to Non-Constant Failure Rates: The Bathtub Curve

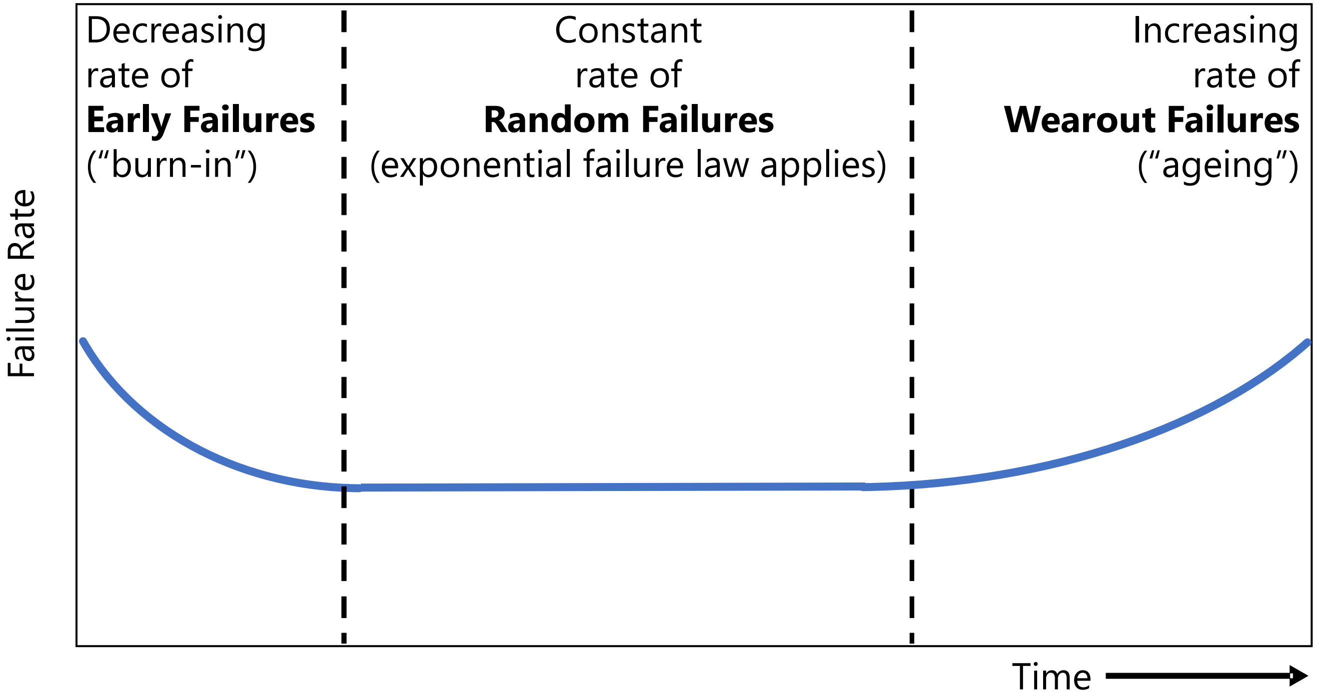

The bathtub curve has three identifiable zones:

1. An early failures zone, the period immediately after manufacture in which there is a relatively higher probability of failure ("infant mortality") which decreases over time
2. A zone of constant and relatively low failure probability, where the exponential failure law applies; this is also called the "random failure plateau"
3. A wearout failures zone, in which the probability of failure begins to increase rapidly with age

In terms of cost growth, the critical segment of the bathtub curve is the "wearout" zone, where age-related effects begin to manifest.

The age of a system significantly influences costs and poses challenges for decision-makers. Age effects are a crucial consideration in determining the cost-effectiveness of replacing an older system. Ultimately, operators aim to establish budgets that are sufficient yet not excessive for maintaining their fleets, such as aircraft, ships, and land vehicles.

This is a preconfigured bathtub curve with an annualized failure rate

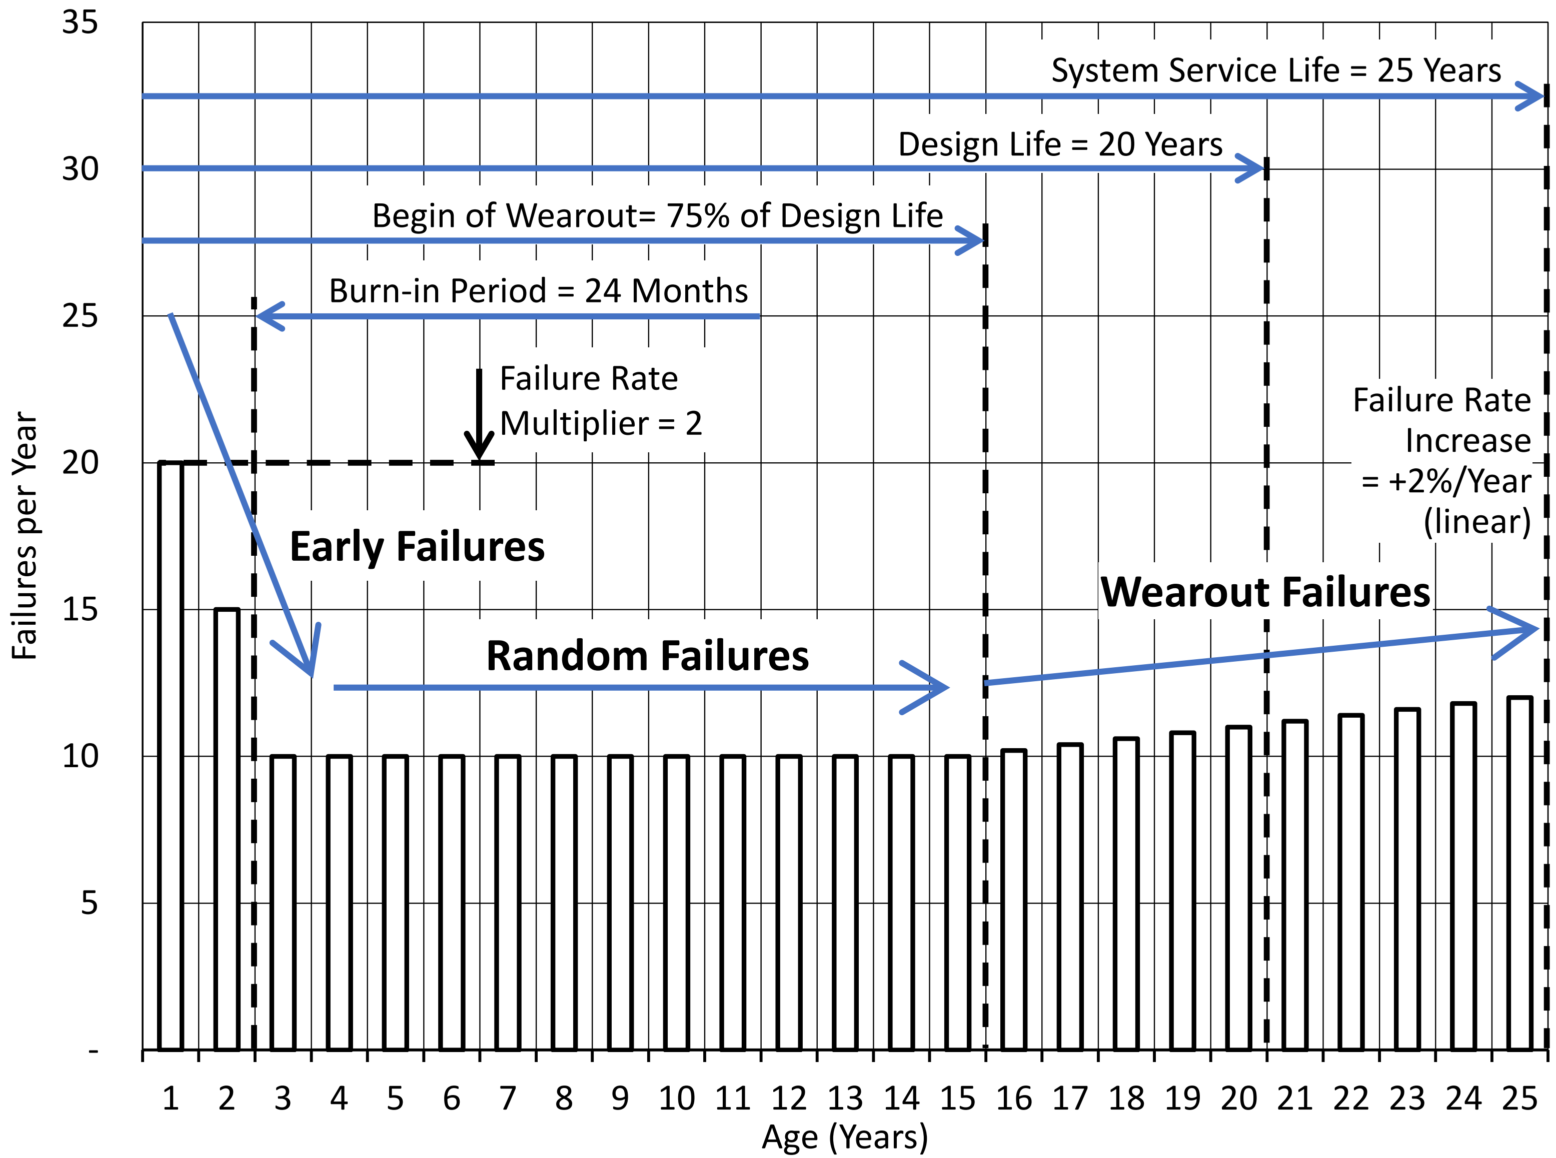

The default configuration of the bathtub curve has an overall failure rate over a 25-year service life that is 10% above purely random failures.

Initially, the curve shows early failures occurring at twice the rate of random failures, meaning the total failure rate at the start of operation is double that of random failures. The “burn-in period” is typically set to 24 months (two years).

The onset of age-related wear is programmed to begin after 75% of the system’s planned design life have elapsed. So, with a design life of 20 years, signs of wear begin to appear after 15 years, manifesting as a linear increase in failure rate by 2% per year (can be switched to exponential).

Users have the flexibility to modify all input parameters as needed.

Additional input parameters for modelling the bathtub curve

Cost model, traditional	Cost model, LCCM
- Total number of systems in use [number, calendar year] - Operating hours [h] - MTBF [h] → **Only random failures with a constant failure rate** can be modelled	- **Number of systems deployed by age cohort** [quantity, calendar year] - System service life [years] *Early failures:* - Burn-in period [months] - Failure rate multiplier at the beginning *Wearout failures:* - System design life [years] - Start of wearout [% of planned design life] - Annual failure rate increase [%] - Annual failure rate increase type [linear or exponential]

One key innovation for LCCM is the way the cost model handles fleet size by introducing age cohorts!

Systems of different ages can show different failure rates within the same calendar year, unless all systems are simultaneously in the section of their bathtub curves with constant failures.

This means that, unlike in traditional models, the number of systems in use and the total operating hours alone are not sufficient to calculate the total failures in a given calendar year. Hence, fleet size must be entered by age cohort.

Total annual failure rate is calculated by stacking bathtub curves of all age cohorts

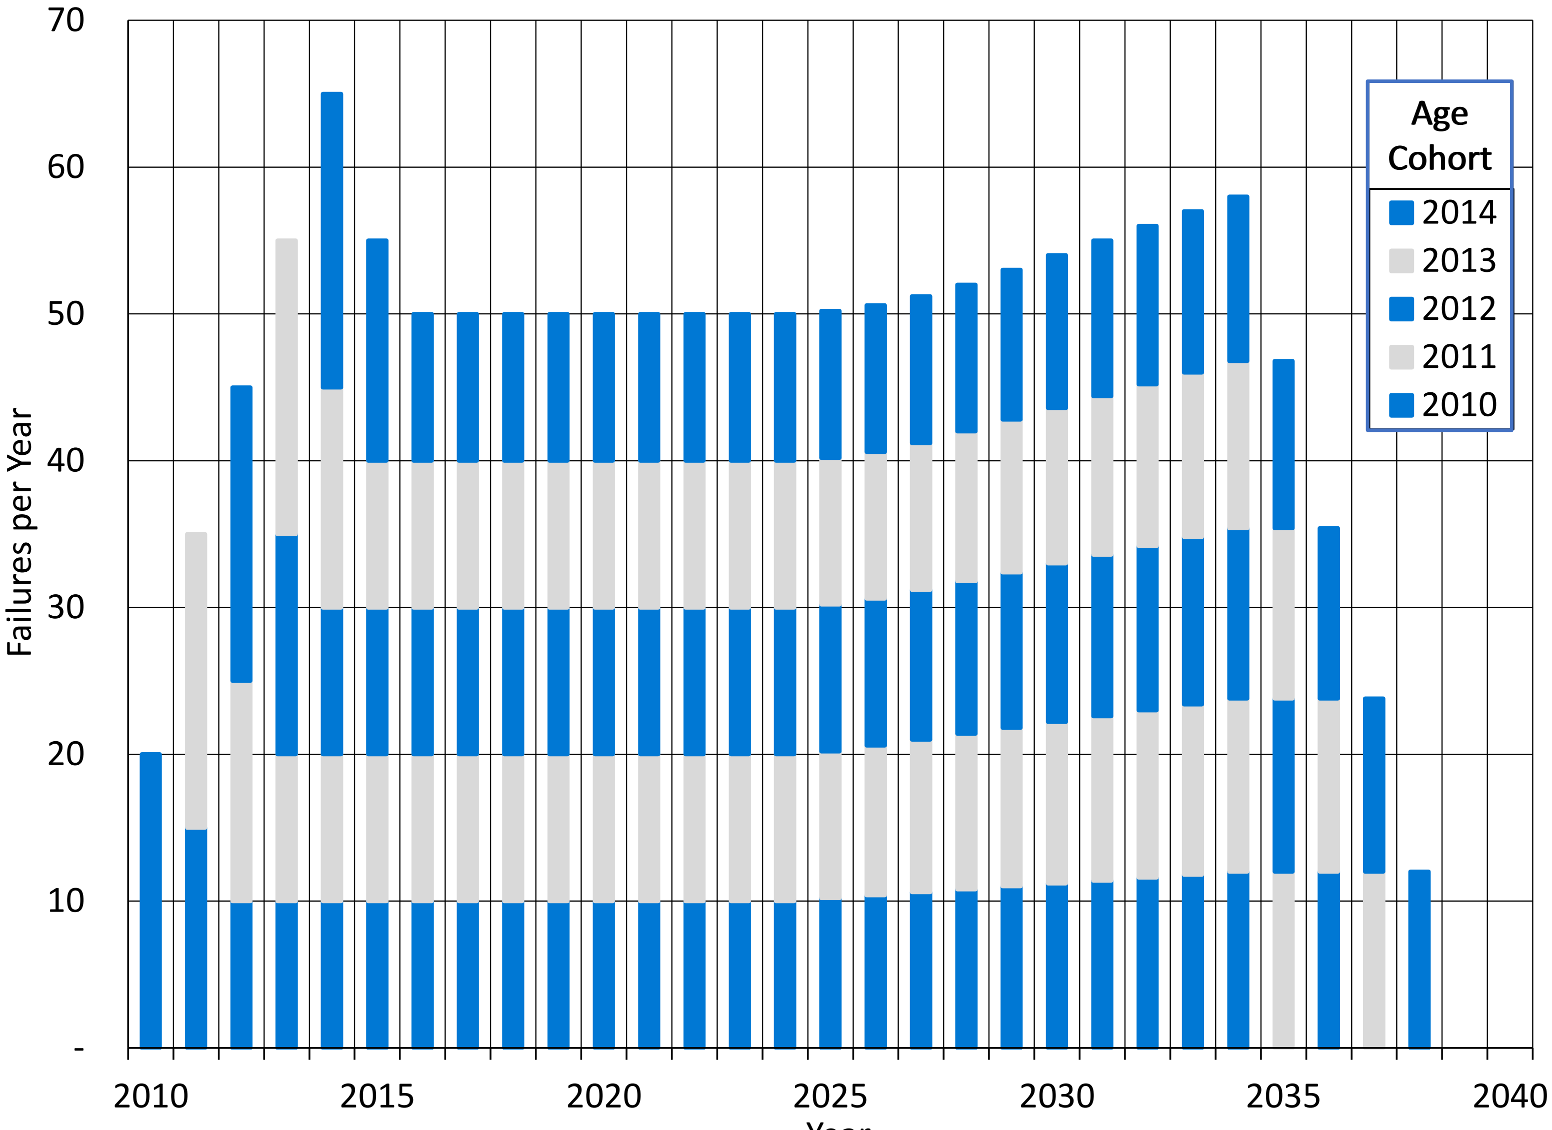

To calculate the total failure rate, each age cohort is assigned its own bathtub curve, with its commissioning year as the starting year.

The total failure rate over time is calculated by adding up the bathtub curves of all age cohorts. This simple example comprises five cohorts of equal size from consecutive calendar years.

The calculation of the total failure rate is shown graphically as a “stacking” of the total of five bathtub curves of all cohorts.

The more age cohorts are involved, and the longer the service life, the more complex shapes of compounded bathtub curves can be created.

Just adding non-constant failures is already a visible departure from the traditional flat Earth model

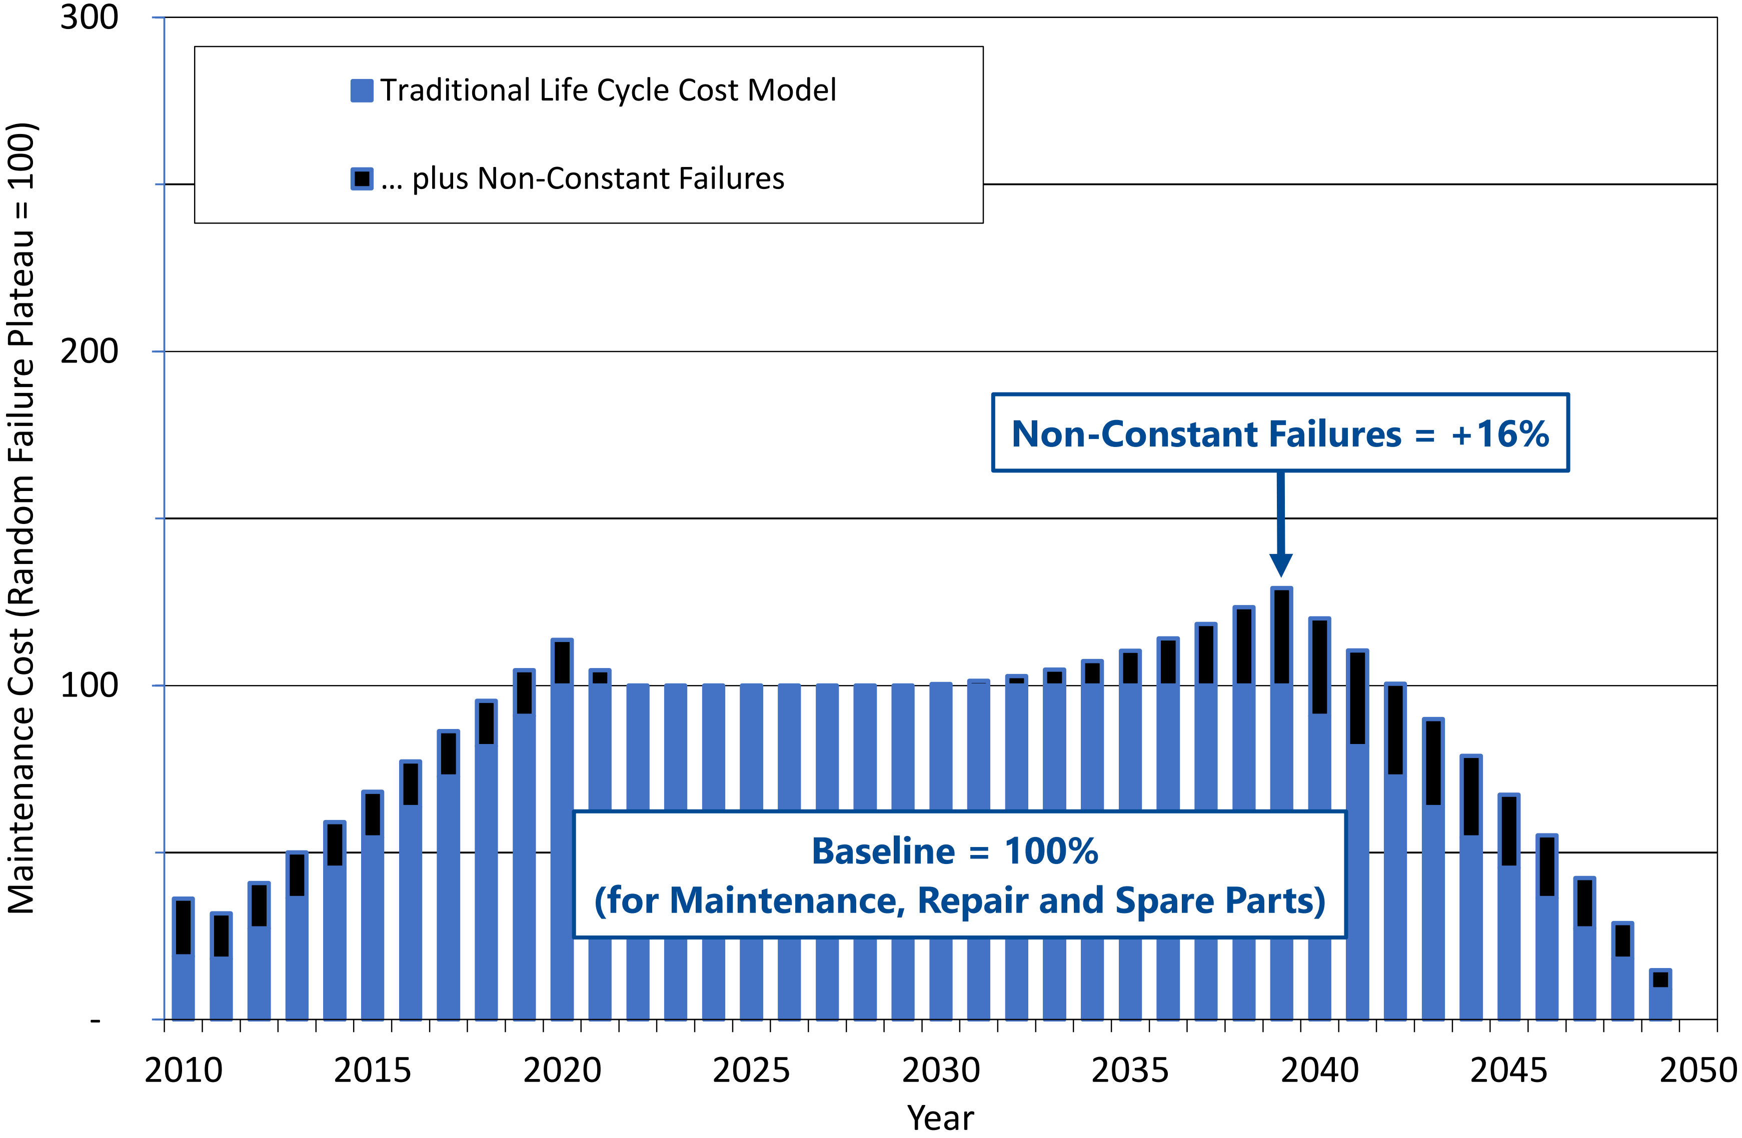

Non-Constant Failures = +16%

Baseline = 100%
(for Maintenance, Repair and Spare Parts)

LCCM advises incorporating the bathtub curve from the outset of a project. This approach allows for an early estimation of the probable increase in operation and support (O&S) costs and their temporal distribution.

LCCM cost models are enhanced with a straightforward, preconfigured bathtub curve to simulate the impacts of variable failure rates. The necessary inputs are easy to modify, **eliminating the need for complex methods such as Weibull distributions.**

Consequently, cost estimators can analyze the financial implications of early and wearout failures **without requiring deep expertise in reliability and failure probabilities.**

2

Condition-Based Maintenance

“No Fault Found”:
The cost of condition-based maintenance

Comparison between corrective and condition-based maintenance

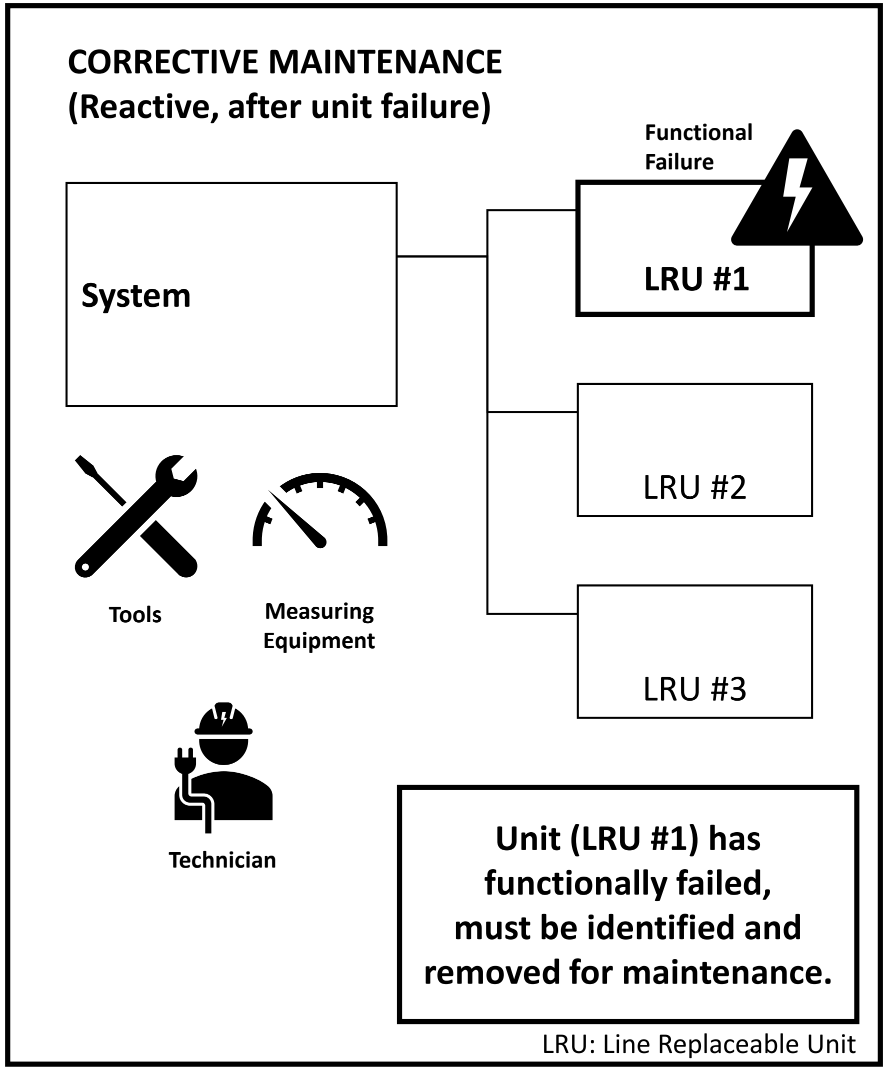

CORRECTIVE MAINTENANCE
(Reactive, after unit failure)

Unit (LRU #1) has functionally failed, must be identified and removed for maintenance.

LRU: Line Replaceable Unit

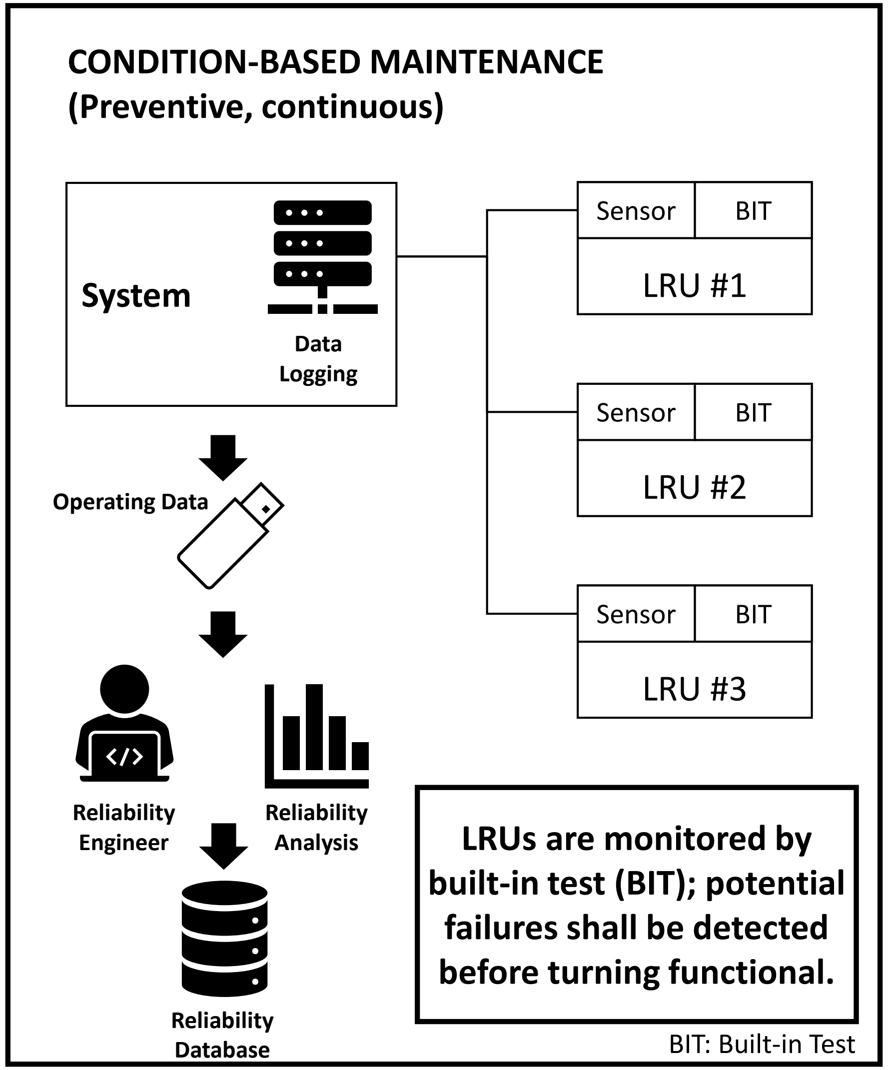

CONDITION-BASED MAINTENANCE
(Preventive, continuous)

LRUs are monitored by built-in test (BIT); potential failures shall be detected before turning functional.

BIT: Built-in Test

Corrective maintenance aims to rectify functional failures only after they have occurred.

Condition-based maintenance aims to improve the efficiency and reliability of systems by identifying potential failures at an early stage, before they can lead to functional failures.

Corrective maintenance is reactive, while condition-based maintenance is preventive.

Workflow of corrective maintenance with 50 faulty units

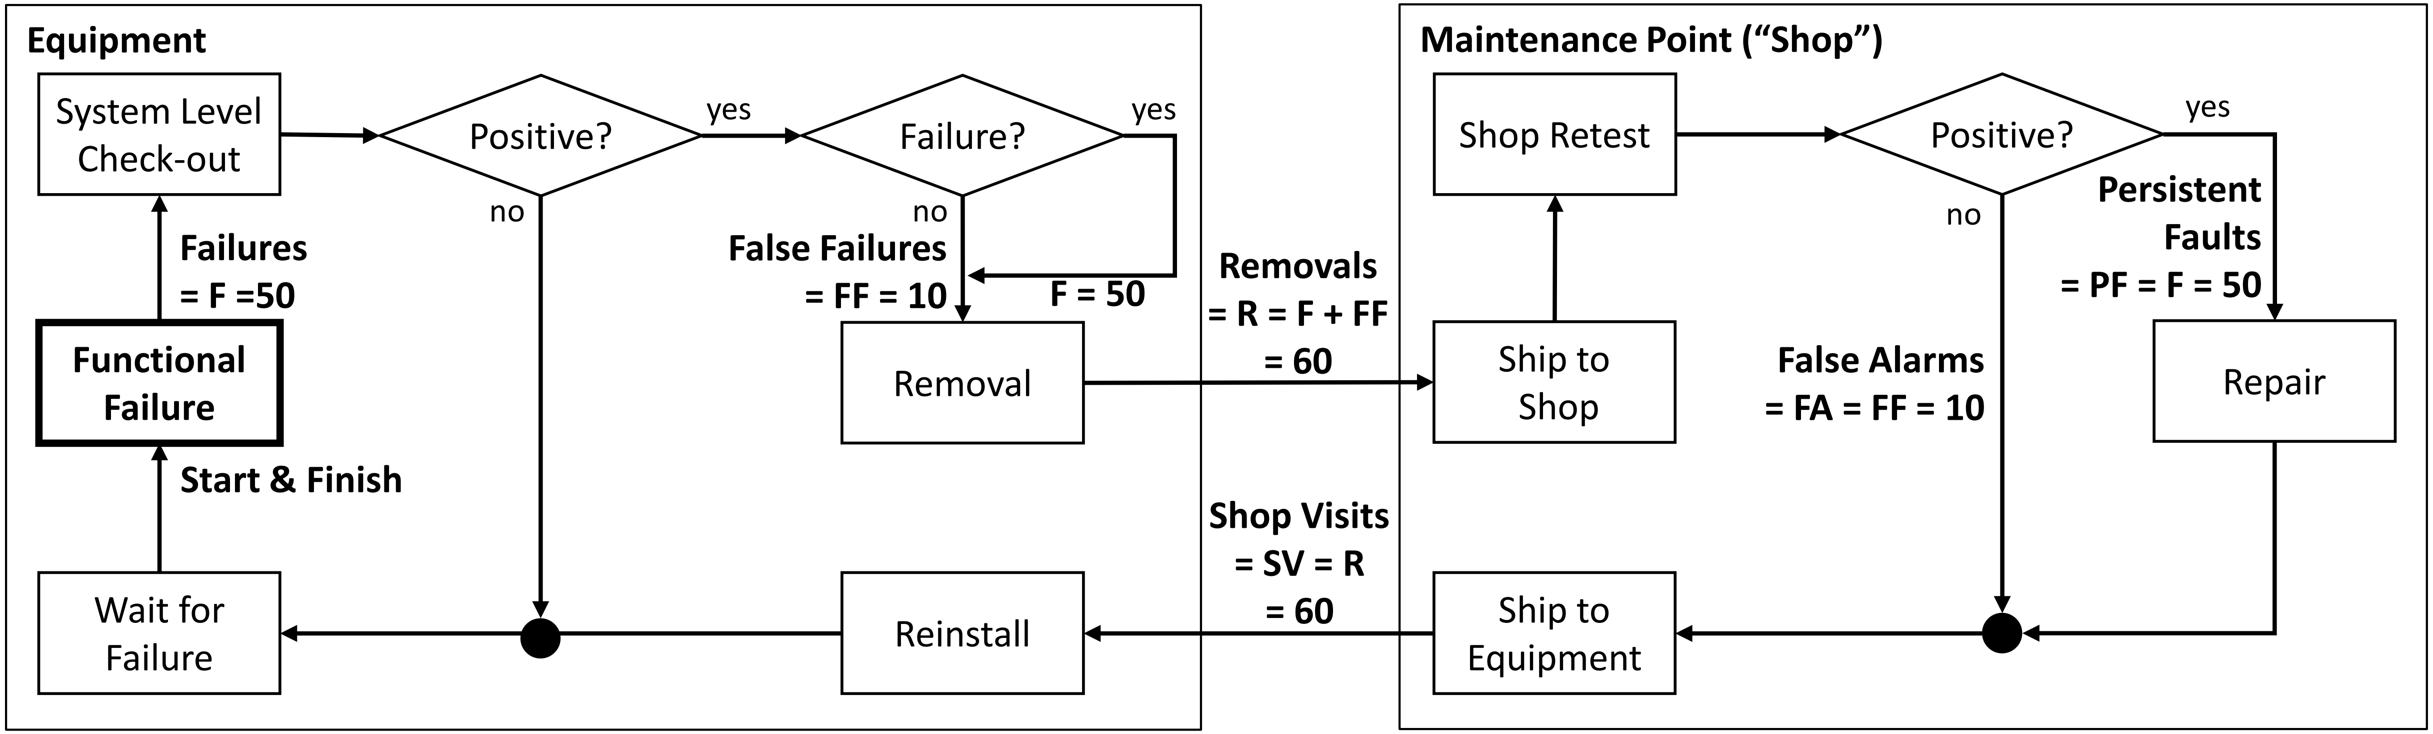

Acronyms: F: Failure; FF: False Failure; FFF: False Failure Fraction; R: Removal; FA: False Alarm; PF: Persistent Fault; SV: Shop Visit

- Number of failures (F): 50; false failure fraction (FFF): 20 %, so number of false failures (FF): 10, because: FF = FFF × F = 0.2 × 50 = 10
- Number of removals (R): 60, because: R = F + FF = 50 + 10 = 60; all 60 removals end up at the maintenance point as a shop visit (SV)
- The 10 false failures (FF) retest negative at the shop, are flagged as false alarms (FA) and sent back unrepaired
- The 50 failed units (F) are confirmed as persistent faults (PF) and repaired

Note: In traditional cost models, fleet size is not factored into the calculation of removals, rendering it irrelevant to the overall cost of maintenance.

Condition-based maintenance means to test, test, test – there are four possible outcomes of a built-in test (BIT) on system level

	Genuinely Positive	**Genuinely Negative**
Test Positive	**True Positive (TP)** • Potential failure detected • Unit is faulty • Unit removed, sent to shop	**False Positive (FP)** • Potential failure detected • But: Unit is good • Unit removed, sent to shop
Test Negative	**False Negative (FN)** • No potential failure detected • But: Unit will functionally fail unless detected in time! • Will need corrective action after unplanned failure!	**True Negative (TN)** • No potential failure detected • Unit is good • Do nothing!

Contrary to corrective maintenance, BIT should proactively identify potential failures before a unit fails and turns dysfunctional.

The effectiveness and costs of condition-based maintenance are critically dependent on the accuracy of fault detection at the system level using built-in tests (BIT).

The BIT must consistently analyse the operating parameters of units to identify any trends that might indicate a potential failure. These trends need to be accurately recognized. In practice, the effectiveness of this process varies, and test results can occasionally be wrong.

Probability of detection (P_{det}) and total fault detection (TFD) are key metrics for the quality of a built-in test (BIT)

The quality of the built-in test (BIT) results depends on the probability of detection (P_{det}), also known as sensitivity. It is the most important parameter of tests at system level. The sensitivity (P_{det}) indicates the probability with which a faulty unit is correctly detected. A sensitivity of 80% means that if a sufficiently large number of tests were carried out and irrespective of the test preconditions, 80% of the units affected by the fault would be detected and 20% of the units affected by the fault would not be recognized as faulty.

$$\textit{Probability of Detection} = P_{det} = \frac{\textit{Faulty units detected}}{\textit{Faulty units tested}} = \textit{Sensitivity}$$

Therefore, 20% of the faulty units tested (and not 20% of tested units) would be false negatives. False negative units that are not detected as faulty in time will suffer unplanned failure during operation and require immediate maintenance. Repeating tests at intervals within the observation period will increase the total fault detection (TFD).

$$\textit{Total Fault Detection} = TFD = 1 - (1 - P_{det})^{n} = 1 - (1 - P_{det})^{\frac{\textit{Observation Period}}{\textit{Inspection Interval}}}$$

Even with a relatively low sensitivity of P_{det} = 80%, the total fault detection rate of a built-in test (BIT) can be effectively increased to 96% by conducting a second test pass, see table below.

Number of tests *n*	**1**	**2**	**3**	**4**	**5**	**>5**
TFD = 1 - (1 - P_{det})^{n} ; P_{det} = 80%	80%	96%	99.2%	99.84%	99.96%	≈100%

Total fault detection of a test with sensitivity P_{det} = 80% and two passes in the observation period is 96%

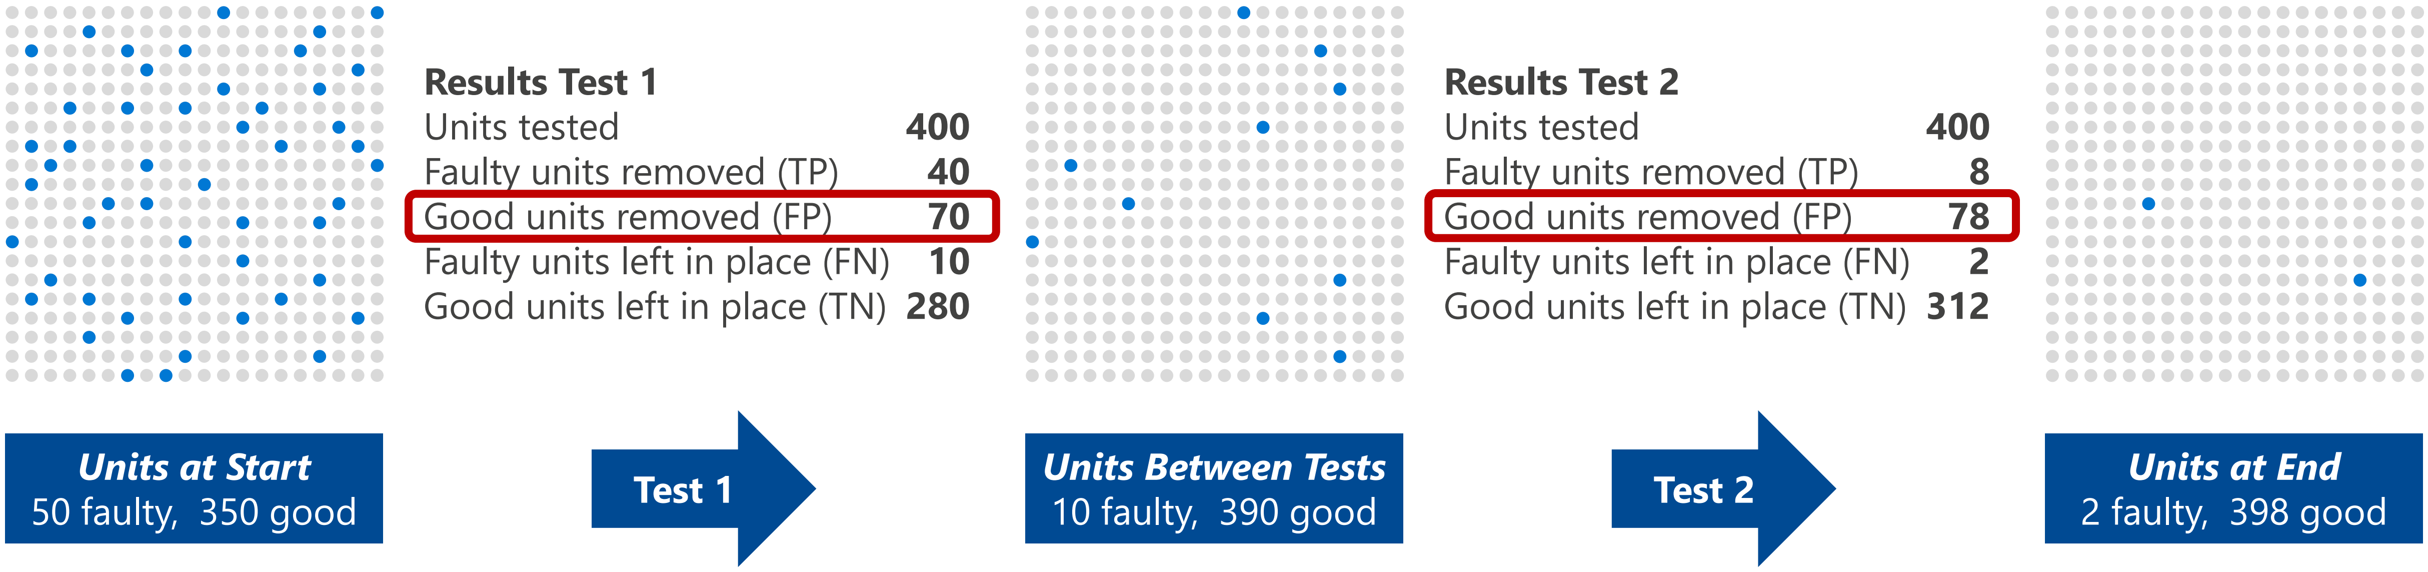

Results Test 1

Units tested	**400**
Faulty units removed (TP)	**40**
Good units removed (FP)	**70**
Faulty units left in place (FN)	**10**
Good units left in place (TN)	**280**

Results Test 2

Units tested	**400**
Faulty units removed (TP)	**8**
Good units removed (FP)	**78**
Faulty units left in place (FN)	**2**
Good units left in place (TN)	**312**

Units at Start
50 faulty, 350 good

Test 1

Units Between Tests
10 faulty, 390 good

Test 2

Units at End
2 faulty, 398 good

TP: True Positive; FP: False Positive; FN: False Negative; TN: True Negative

System-Level Test outcomes after an observation period with two passes of a BIT **with P_{det} = 80%**:

- **Probability of detection P_{det} = 80%; so, → False Positive Rate (FPR) = (1 - P_{det}) = 20%**
- Total number of units tested: 400 + 400 = 800
- Total faulty units removed: 0.8×50 + 0.8×10 = 40 + 8 = 48 True Positives (TP) → **Total Fault Detection = 48/50 = 96% → Good!**
- Total good units removed: 0.2×350 + 0.2×390 = 70 + 78 = 148 False Positives (FP) → **False Positive Rate = 148/(350 + 390) = 20% → Bad!**
- Total of undetected faulty units suffering **Functional Failure** by the end of the observation period: 50 - (40 + 8) = 2 False Negatives (FN)
- The total number of units removed after a positive BIT or functional failure: TP + FP + FN = 48 + 148 + 2 = **198 Units Under Test (UUT)**

→ Observation: The number of False Positives increases with the number of inspections during the observation period!

What happens after removal from the system and shipping to the shop? There are three different outcomes of a shop retest

	Genuinely Positive	Genuinely Negative
Retest Positive	**Persistent Fault (PF)** • Unit fails retest, is confirmed as faulty • Unit is repaired, sent back to system • Unit is reinstalled, passes check-in at system	(Not applicable)
Retest Negative	**Intermittent Failure (IF)** • Unit passes retest as good • Unit is sent back to system unrepaired • Unit is reinstalled, fails check-in at system • Unit is removed, sent to shop again …	**False Alarm (FA)** • Unit passes retest as good • Unit is sent back to system unrepaired • Unit is reinstalled, passes check-in at system

The cost model must count and assign different cost to the three different outcomes of the Shop Retest and their impact on maintenance cost, like:

- Faulty units that came as (true) positives and are retested positive (Persistent Faults) are good; they are repaired and returned to equipment
- Faulty units that came as (true) positives and are retested negative (Intermittent Failures) are bad; they need additional cycles BIT–Shop Retest until they turn into Persistent faults
- **Good units that came as (false) positives and are retested negative (False Alarms) are very bad; they do not need repair and cause meaningless effort for testing, removal, and shipping**

Workflow for condition-based maintenance with 50 faulty units out of 400 and two test passes at system level

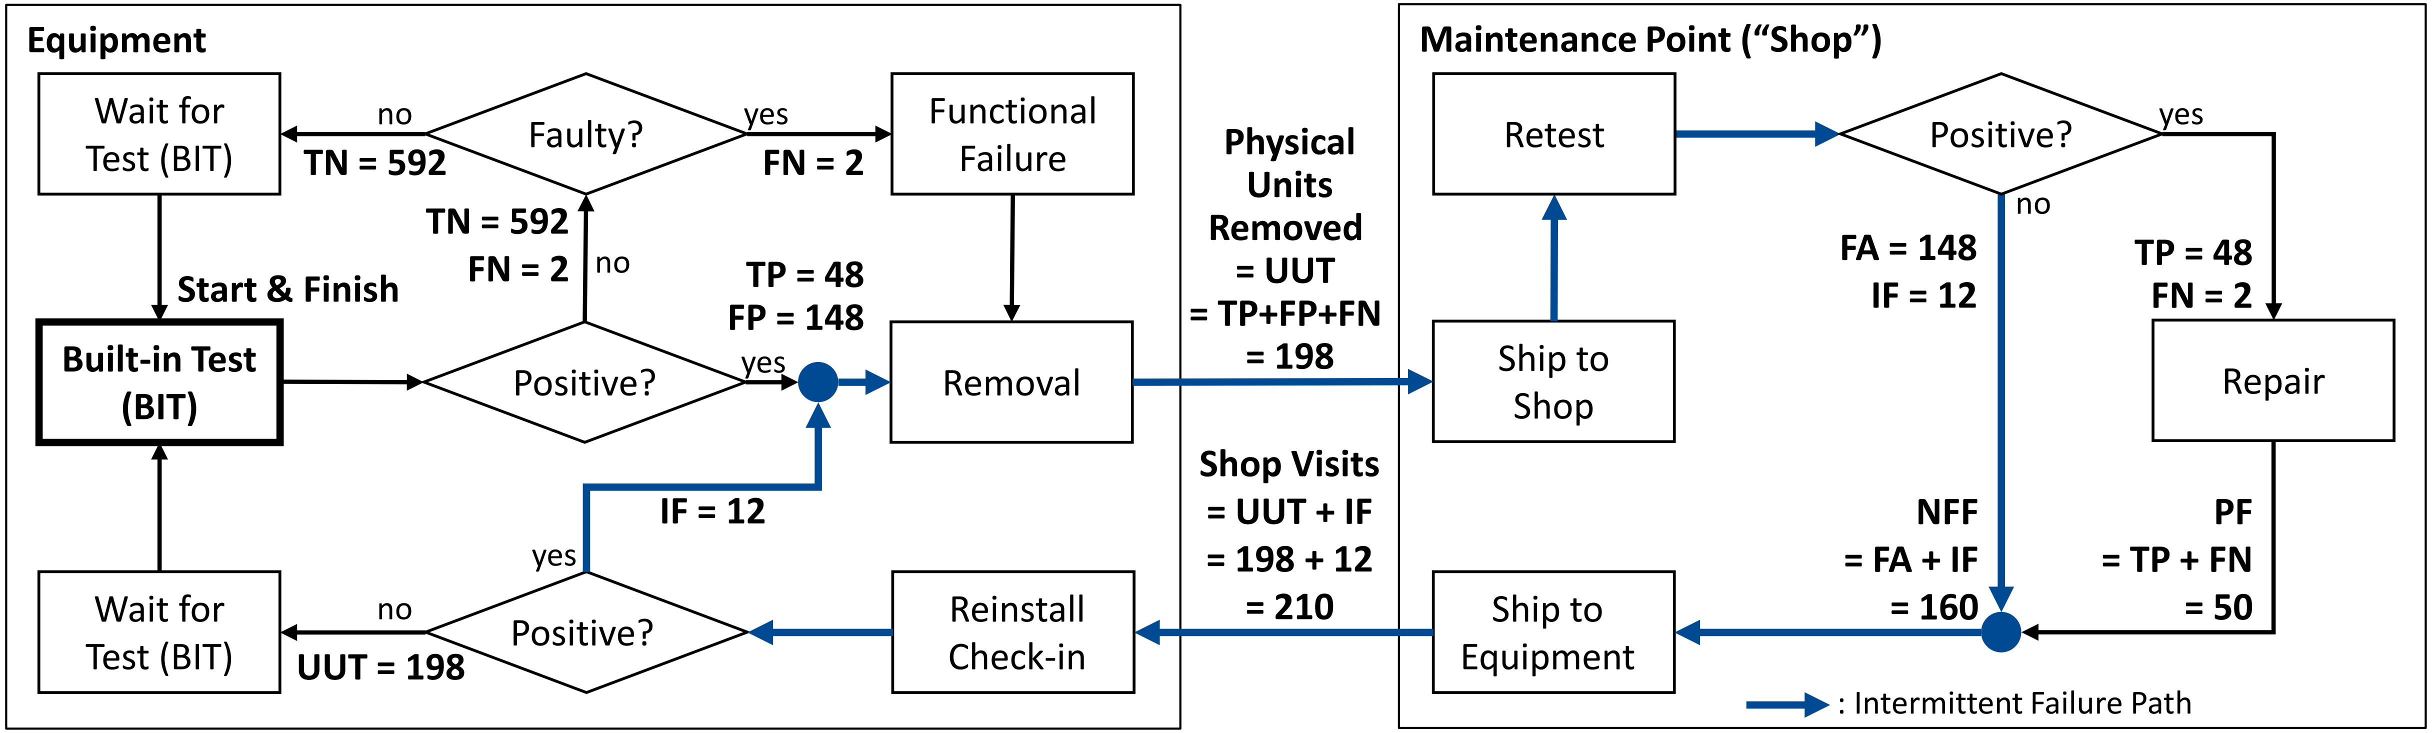

TP: True Positive; FN: False Negative; FP: False Positive; TN: True Negative; UUT: Units Under Test; PF: Persistent Fault; IF: Intermittent Failure; FA: False Alarm; NFF: No Fault Found; SV: Shop Visit. Note: To arrive at the cumulative total of 800 tests (BIT) after the second pass, FN = 10 from the first pass must be added!

- With 50 faulty units, the total number of units removed and shipped to the shop is UUT = TP + FP + FN = 48 + 148 + 2 = 198
- With an intermittent failure rate (IFR) of 20%, it takes the 48 true positives an extra 0.2/(1 - 0.2) × 48 = 12 shop visits = Intermittent Failures (IF)
- Total shop visits SV = UUT + IF = 198 + 12 = 210
- No Fault Found NFF = SV - PF = 160 = UUT + IF. **During 160 out of 210 shop visits, the result is No Fault Found (NFF)!**

Comparing false alarms and no fault found makes condition-based maintenance look extremely bad

Maintenance Type	F, PF	UUT	IF	FA	NFF	SV	NFF%
Corrective	50	60	-	10	10	60	**16.67%**
Condition-based	50	198	12	148	160	210	**76.19%**

For condition-based maintenance of 50 faulty units (P, PF = 50):

- … the number of units under test (UUT) removed at system level is more than three times higher (UUT = 198 instead of 60)
- … the number of No Fault Found (NFF = IF + FA) is 16 times higher (NFF = 160 instead of 10)
- … the number of shop visits SV (SV = UUT + IF) is increased by a factor of 3.5 (SV = 210 instead of 60)
- … the fraction of shop visits with “no fault found” (NFF% = NFF/SV) is increased by a factor of more than 4.5 (NFF% = 76.19 % instead of 16.67 %)

This mainly comes from the nearly 15-fold (!) increase in the number of false alarms (FA = 148 instead of 10).
Intermittent failures (IF) only occur in condition-based maintenance, but their overall impact is negligible (IF = 12 instead of 0)

These values may seem exaggerated, but they are not. Real-world examples show NFF of up to 90% and more for condition-based maintenance, with the majority being due to false alarms. The costs of this issue are considerable. In 2012, the US Department of Defense announced that it was spending two billion dollars a year (!) on the removal of line replaceable units (LRUs) that were tested with NFF during the incoming inspection at the depot. Experts say that they will never be able to eliminate all NFF; however, they are confident the extent of the problem can be reduced *(Werner, D. A maddening, costly problem. Aerosp. Am. 53.2, pp. 28–33 (2015)).*

The major cause of No Fault Found are False Positives leading to False Alarms

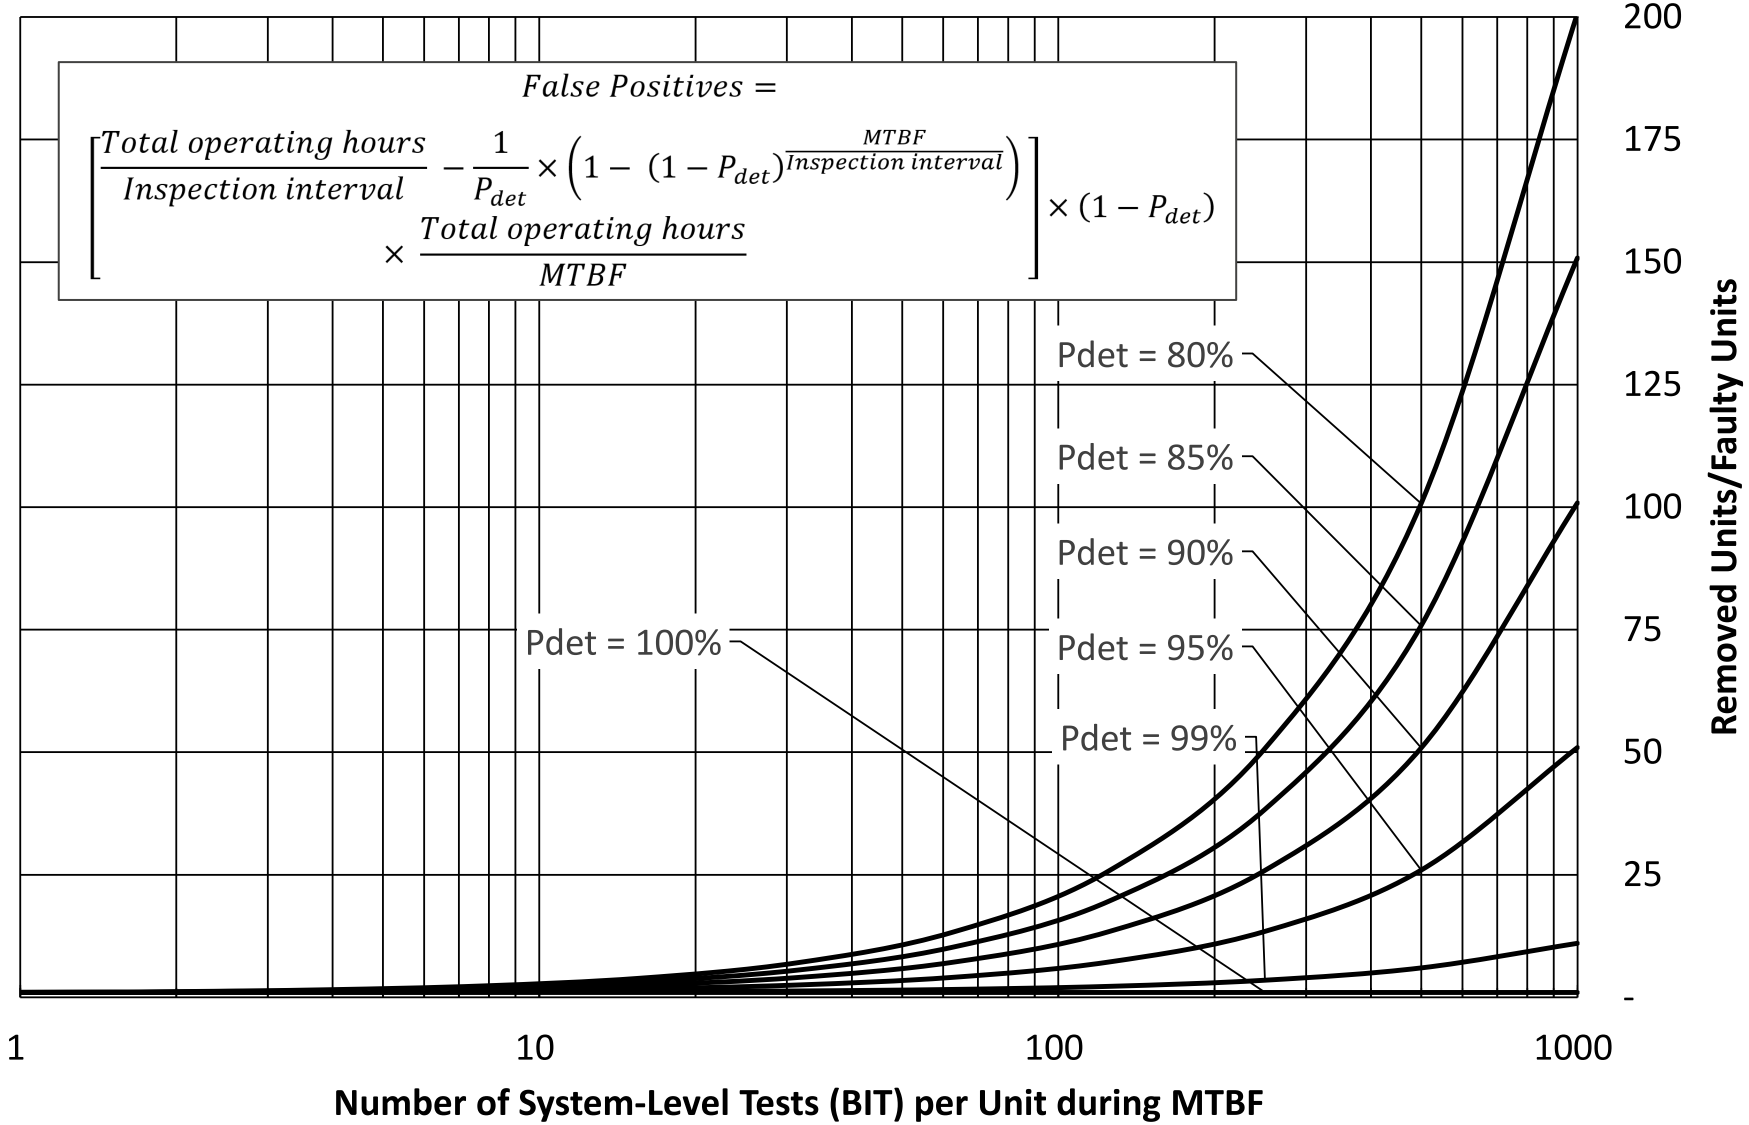

$$\text{False Positives} = \left[\frac{\text{Total operating hours}}{\text{Inspection interval}} - \frac{1}{P_{det}} \times \left(1 - (1-P_{det})^{\frac{MTBF}{\text{Inspection interval}}}\right) \times \frac{\text{Total operating hours}}{MTBF}\right] \times (1-P_{det})$$

The disadvantage of frequent testing is that faulty units are recognised and removed with each additional test. So, the number of genuine negatives increases with every pass. The false positive rate (1 - P_{det}), however, remains constant.

The lower the sensitivity (P_{det}) of the test at system level and the shorter the inspection interval, the greater the number of false positives removed!

For the sake of affordability, such extreme conditions will not be acceptable in practical applications. It is more likely that maintenance will ensure that P_{det} at system level is close to 100%.

Additional input parameters for modelling condition-based maintenance

Cost model, traditional	Cost model, LCCM
- Operating hours [h] - MTBF (Mean Time Between Failures) [h] - MTTR (Mean Time to Repair) [h] → **Only corrective maintenance after random failures with a constant failure rate** can be modelled	- Probability of detection P_{det} at start [%] - Probability of detection P_{det} at goal [%] - Goal, in cumulative time of operation [hours] - Inspection interval [hours] - Intermittent failure rate [%] - Duration BIT at system level [hours] - Continuous incoming inspection for maintenance [hours] - Duration of function test after reinstallation [hours] - Unplanned failure penalty in labour hours [hours] - Unplanned failure penalty in material cost [€]

The probability of detection (P_{det}) by the built-in test (BIT) is the most critical cost driver. Given its significance, **P_{det} must be treated as a variable within the LCCM cost model**. This enables a detailed analysis of learning effects in error detection and their impact on costs. The model requires the following inputs:

- **An initial sensitivity value (P_{det}) at the start of the operation and support (O&S) phase (in %)**
- **A milestone defined by the cumulative operating time of the entire fleet (in hours)**
- **A target sensitivity value (P_{det}) to be achieved by the time of the milestone (in %)**

For instance, at the onset of the O&S phase, the system-level BIT might have an initial detection probability (P_{det}) of 80%. After the fleet accumulates 100,000 operational hours, the target for P_{det} shall be at 95%. This allows for an evaluation of improvements in fault detection capabilities over time.

No Fault Found (NFF) have a big impact on maintenance cost

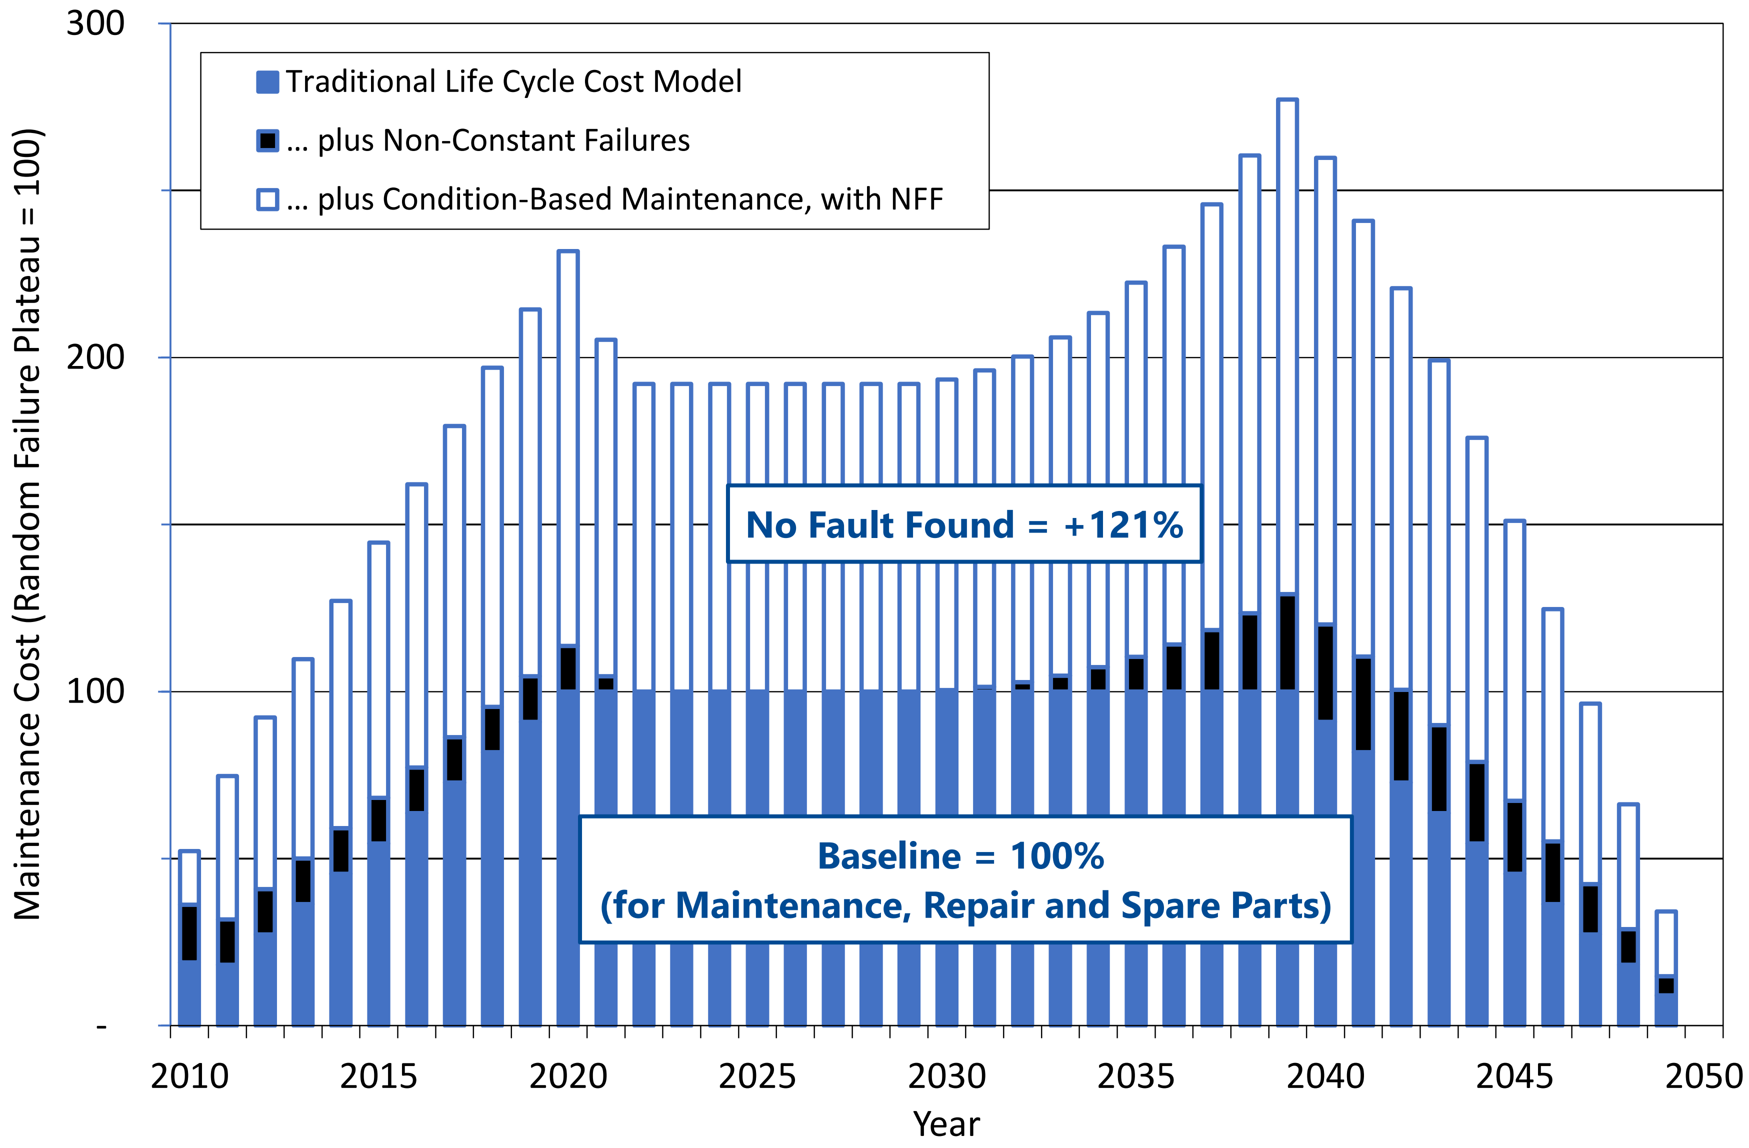

Modern aerospace and defence systems, rich in software and electronics, rely on built-in tests (BIT) for condition-based maintenance (CBM).

While essential, BIT often generate false positives, leading to unnecessary removal and maintenance of perfectly functional units. These are frequently retested as “no fault found” (NFF), wasting time, effort, and resources.

Understanding the cost drivers associated with CBM, especially those related to NFF, is crucial. Improved insight into these costs can help refine maintenance protocols and BIT configurations, reducing inefficiencies and enhancing system reliability and operational efficiency.

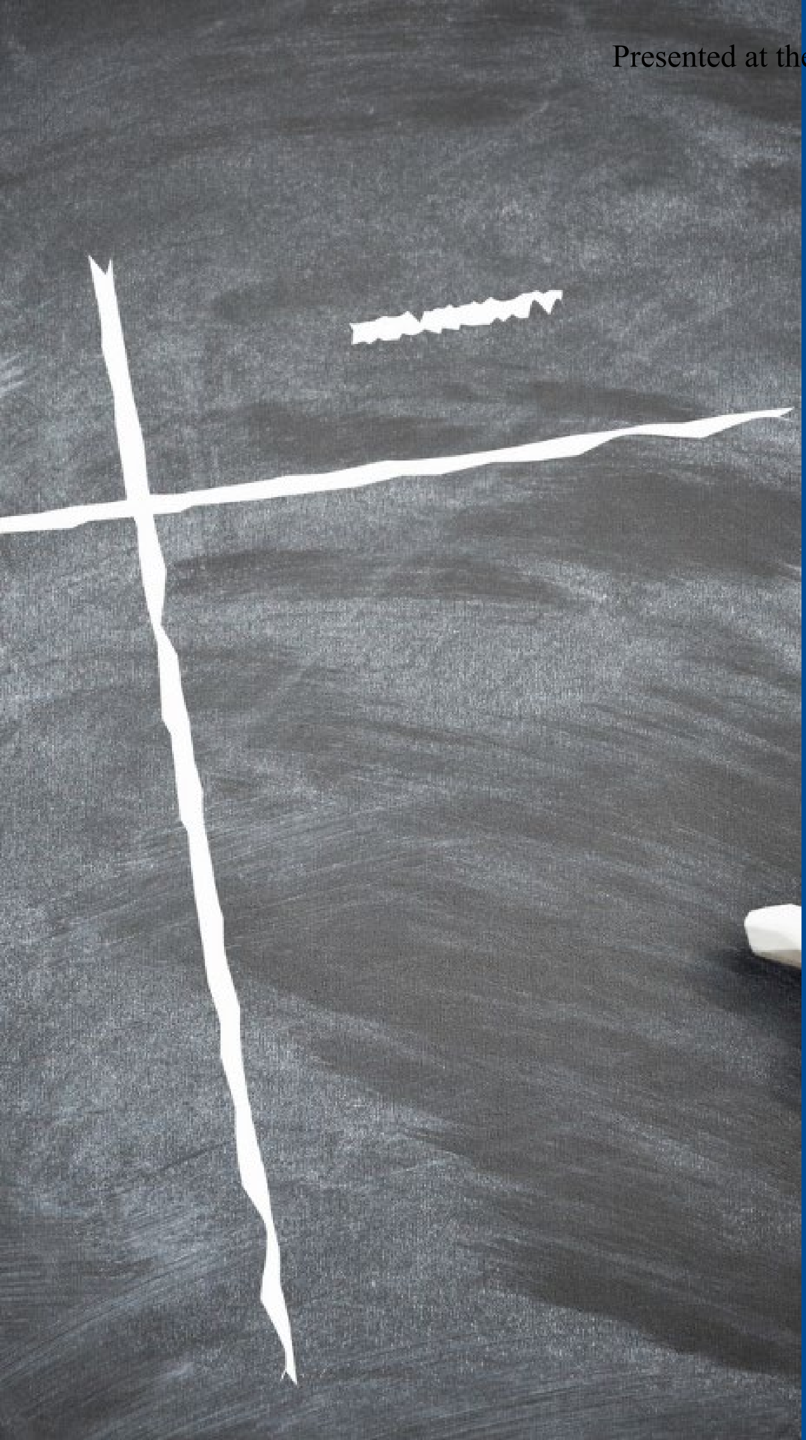

3

Spare Parts Obsolescence

"Parts Obsolete": The cost of product and technology obsolescence

With LCCM, replenishment spares cost is using a new logic

“Flat Earth” Spares Cost Logic

- Spares technology is “frozen” from start of original production phase
- Spares cost is based on the original production quantity
- Supply of spares will be available throughout service life, regardless of duration
- Even if service life goes over 50 years or longer, spares will not change
- Spares cost algorithms date back to 1970s

LCCM Spares Cost Logic

- Obsolescence of spares is real and can be mitigated
- Spares can be replaced by more recent successors, reflecting technology improvements
- Cost adjustments for more recent spares can rely on existing technology improvement and obsolescence control models

There are two kinds of obsolescence – Product Obsolescence and …

Product obsolescence occurs when a product is discontinued by its original manufacturer, often due to a variety of causes. **Product cycles are relatively short, 3 to 7 years are typical values.**

Reasons can include the unavailability of necessary materials, reduced demand for the component, or liability concerns. This issue is particularly acute in the realm of electronics, where the procurement lifetimes of microelectronic components are frequently much shorter than the manufacturing and support lifecycles of the systems that utilize these components.

Technology Obsolescence

Technology obsolescence occurs in long cycles – 40 years are typical – as new products based on advanced technologies enter the market, reducing demand for older technology lines.

This makes it less viable for manufacturers to maintain production capacity for outdated products. Initially, remaining stock can be used or production outsourced, but over time, the number of suppliers dwindles, decreasing availability and expert knowledge. So, prices rise, contrary to the effects of technology improvement.

Eventually, when no suppliers remain for older parts, technology obsolescence is fully realized.

Before estimating obsolescence cost, there were models for technology improvement and the "Z-curve"

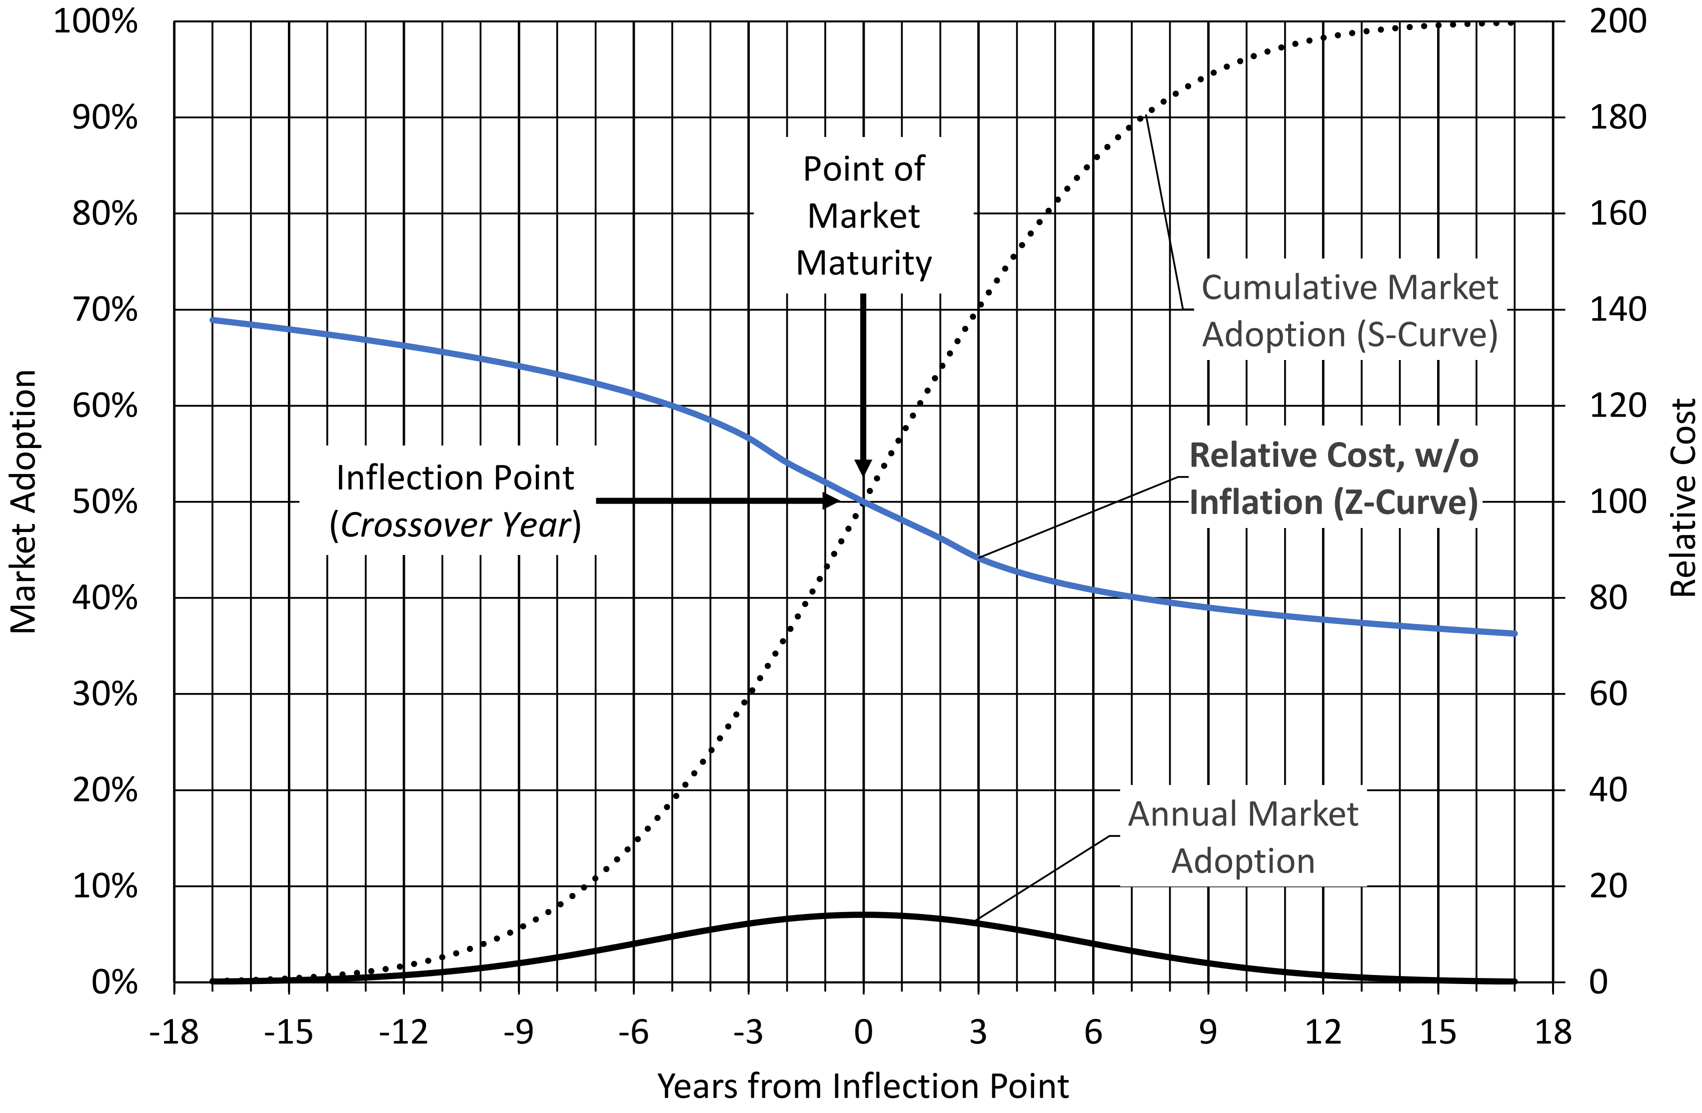

Since around 1975, commercial cost estimation tools (e.g., PRICE) have been using built-in models to **calculate the impact of technology improvement on manufacturing costs**. They are based in part on the innovation diffusion model developed by Everett Rogers in 1962.

While cumulative market penetration follows the characteristic S-curve upwards, adjustment of relative product costs is directed downwards. It resembles an inverted S-curve, which is why it is also referred to as the "Z-curve".

Fun fact: Z-curves adjusted for technology obsolescence were only introduced in 2006.

“Flat Earth” Fiction #1: Spares will never change and last forever

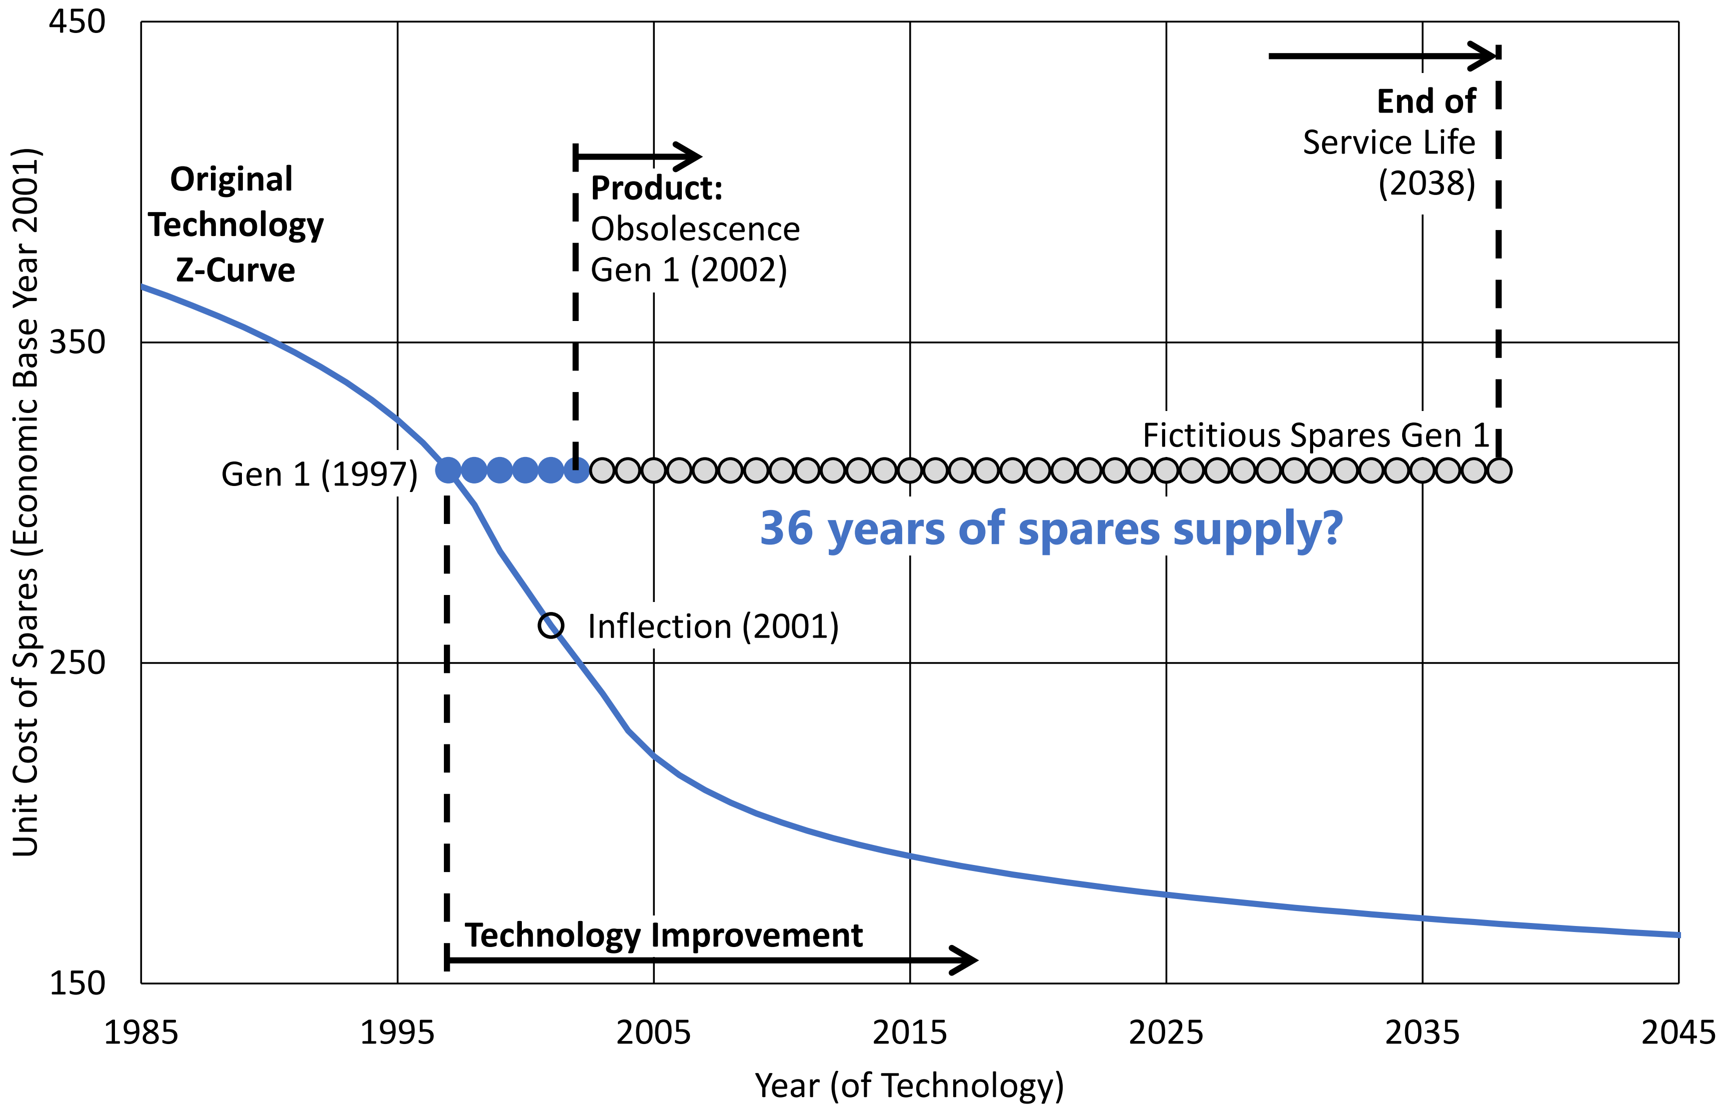

Traditional Cost Model (Example):

- A product is based on a technology following a Z-Curve with the inflection point in 2001
- When the first product generation (Gen 1) enters production in 1997, its technology is frozen
- The product is used in a system whose service life goes until 2038
- Regardless of the 42-year time horizon, the model assumes and calculates that spares will last forever and stay the same
- With a Gen 1 product discontinued after 6 years, the spares supply until 2038 may be called “fictitious”

"Flat Earth" Fiction #2: Traditional cost models can now do technology obsolescence control (sort of)

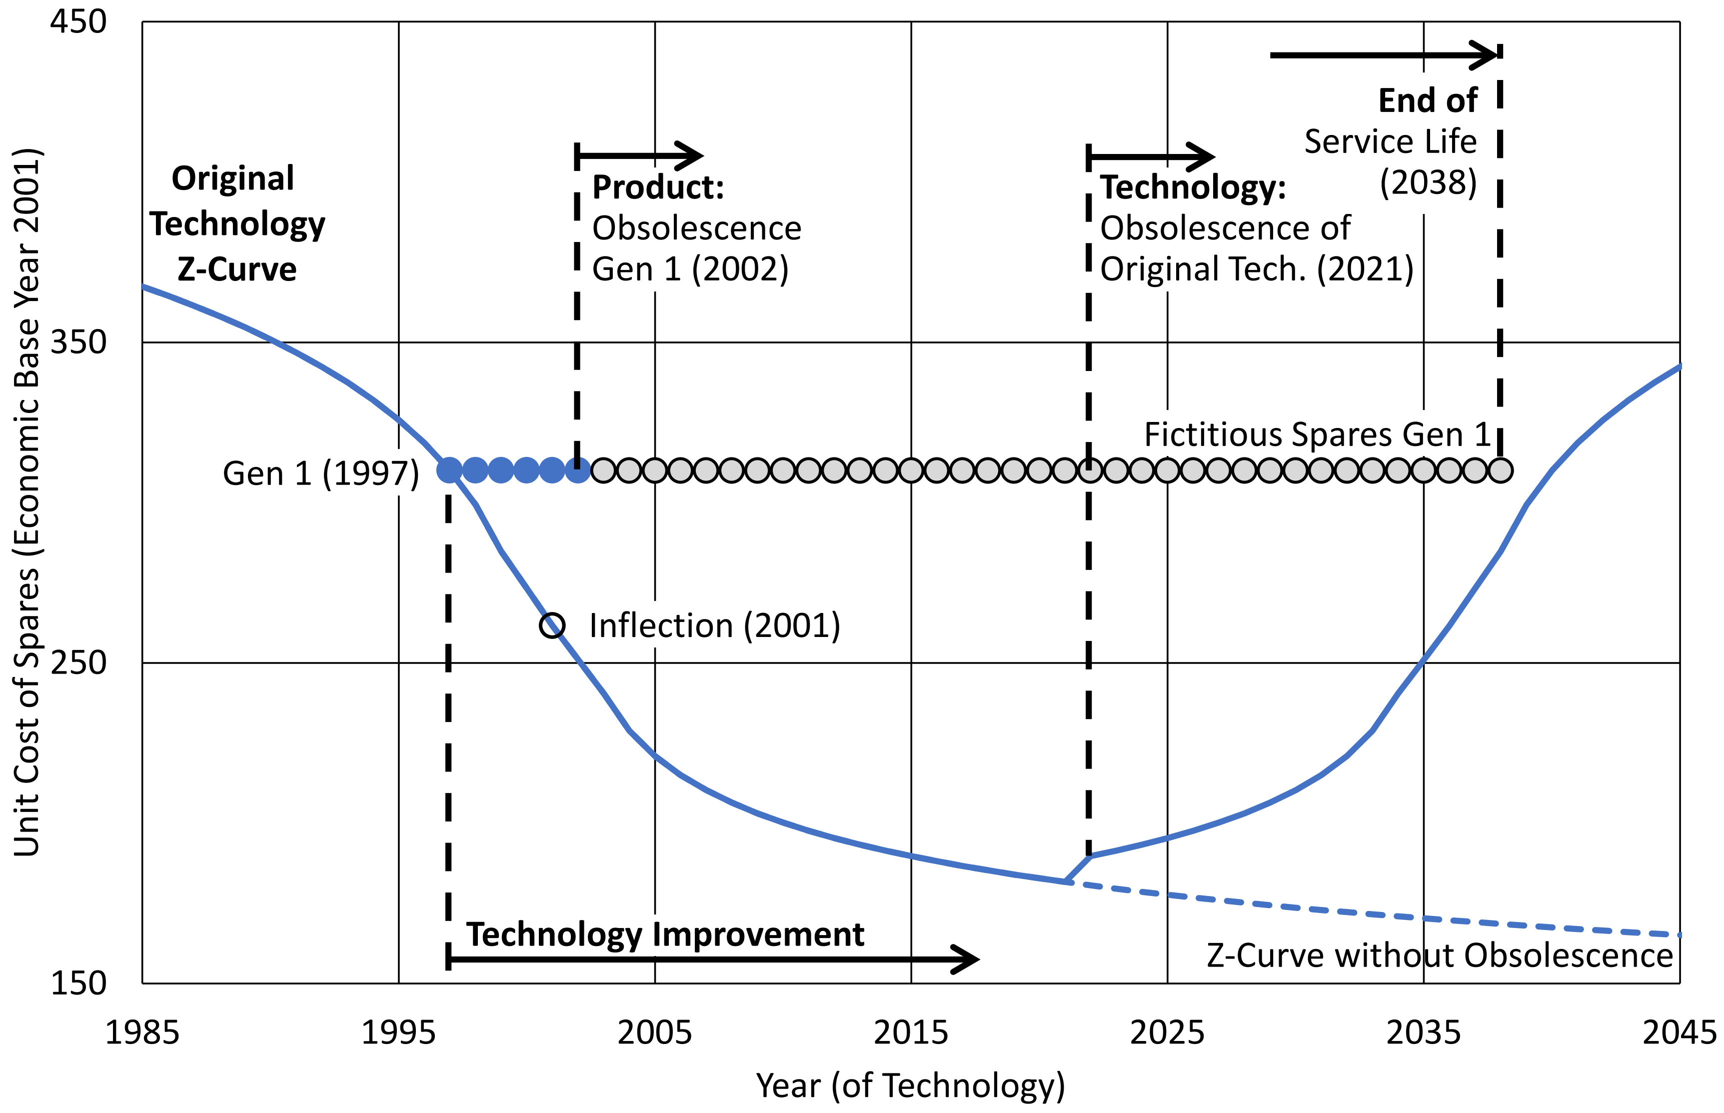

Traditional Cost Model:

- Z-curves can switch on technology obsolescence; by default it is off
- If switched on, technology can still be frozen **only once**, at production start!
- It might be worth checking if a cost estimate contains technology whose obsolescence has already begun
- The inflection point may need to be moved on the time axis (it can be done, but it is scary)
- The technology obsolescence control feature might be too confusing to use
- **The original fiction persists: Technology will stay the same for the whole service life!**

LCCM Fact #1: Spare parts can surf the Z-curve in the same way as the original product, one generation at a time

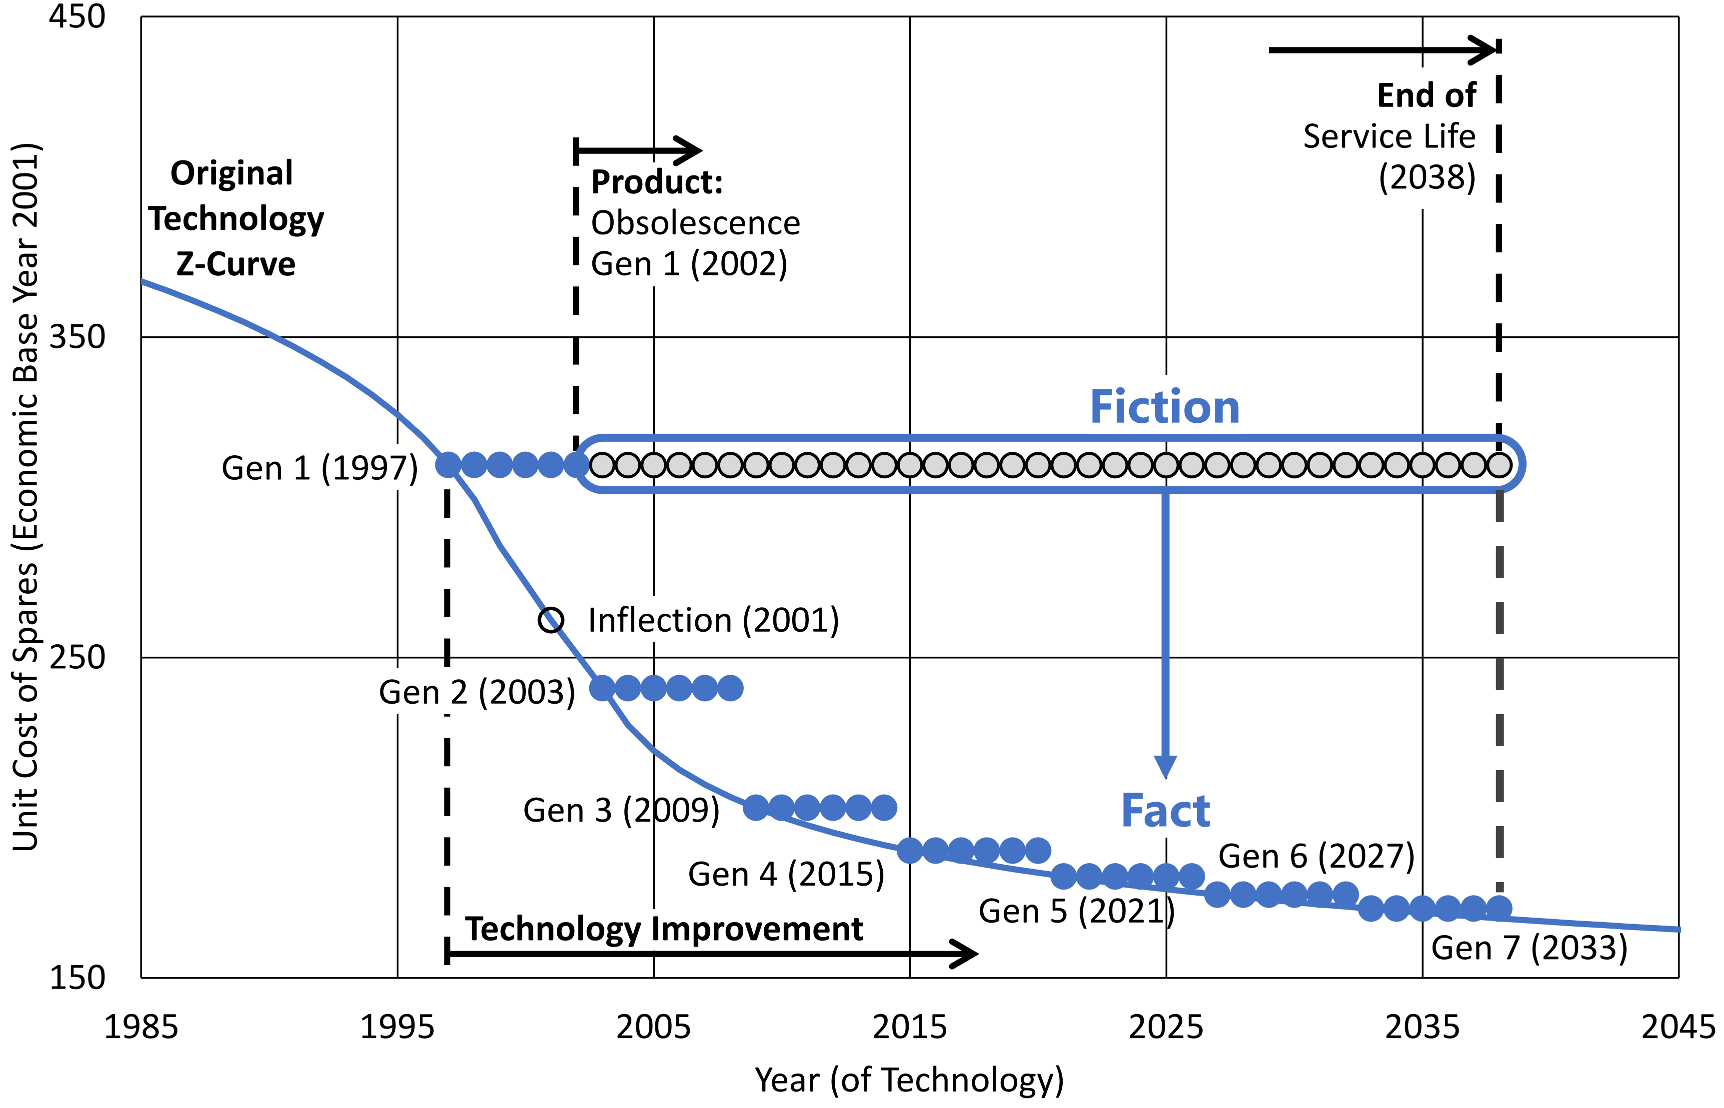

LCCM Cost Model:

- **An LCCM cost model allows to control technology improvement of spare parts independently from the original product**
- The fiction of an endless supply of original spares is replaced by fact – a string of product generations surfing along the Z-curve
- With every new generation, its inherent technology is frozen at a later year of technology, the unit cost of spares is adjusted accordingly

LCCM Fact #2: Finally, it makes sense to have a Z-curve with technology obsolescence adjustment

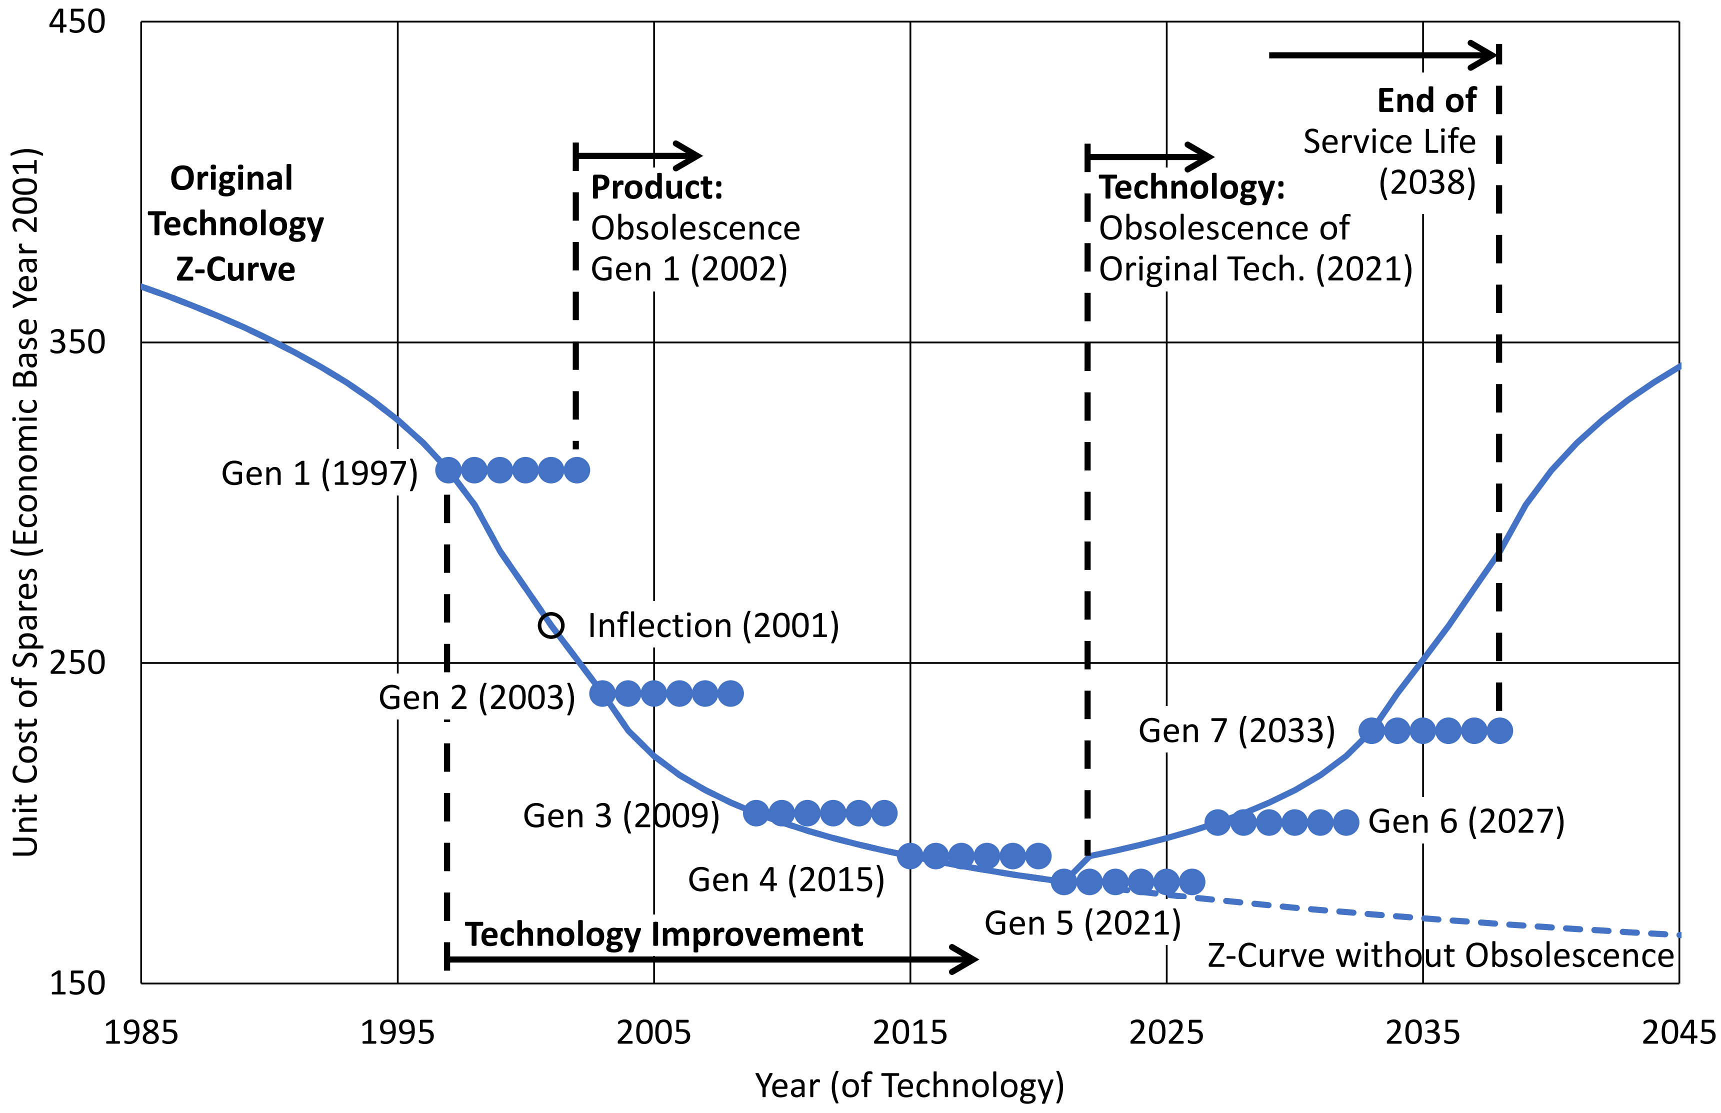

LCCM Cost Model:

- For five generations, Gen 1 to Gen 5, the product family can surf down the Z-curve, seeing 4 instances of **product obsolescence**
- Technology improvement leads to cost savings from one generation to the next
- Gen 5 sees a minimum in unit cost, before **technology obsolescence** kicks in after 2021
- Technology obsolescence leads to cost increases from Gen 6 onwards
- The unit cost of the ultimate Gen 7 (2033) is between Gen 2 (2003) and Gen 3 (2009)

LCCM Fact #3: Now spares can even change to a higher technology when their original technology becomes obsolete

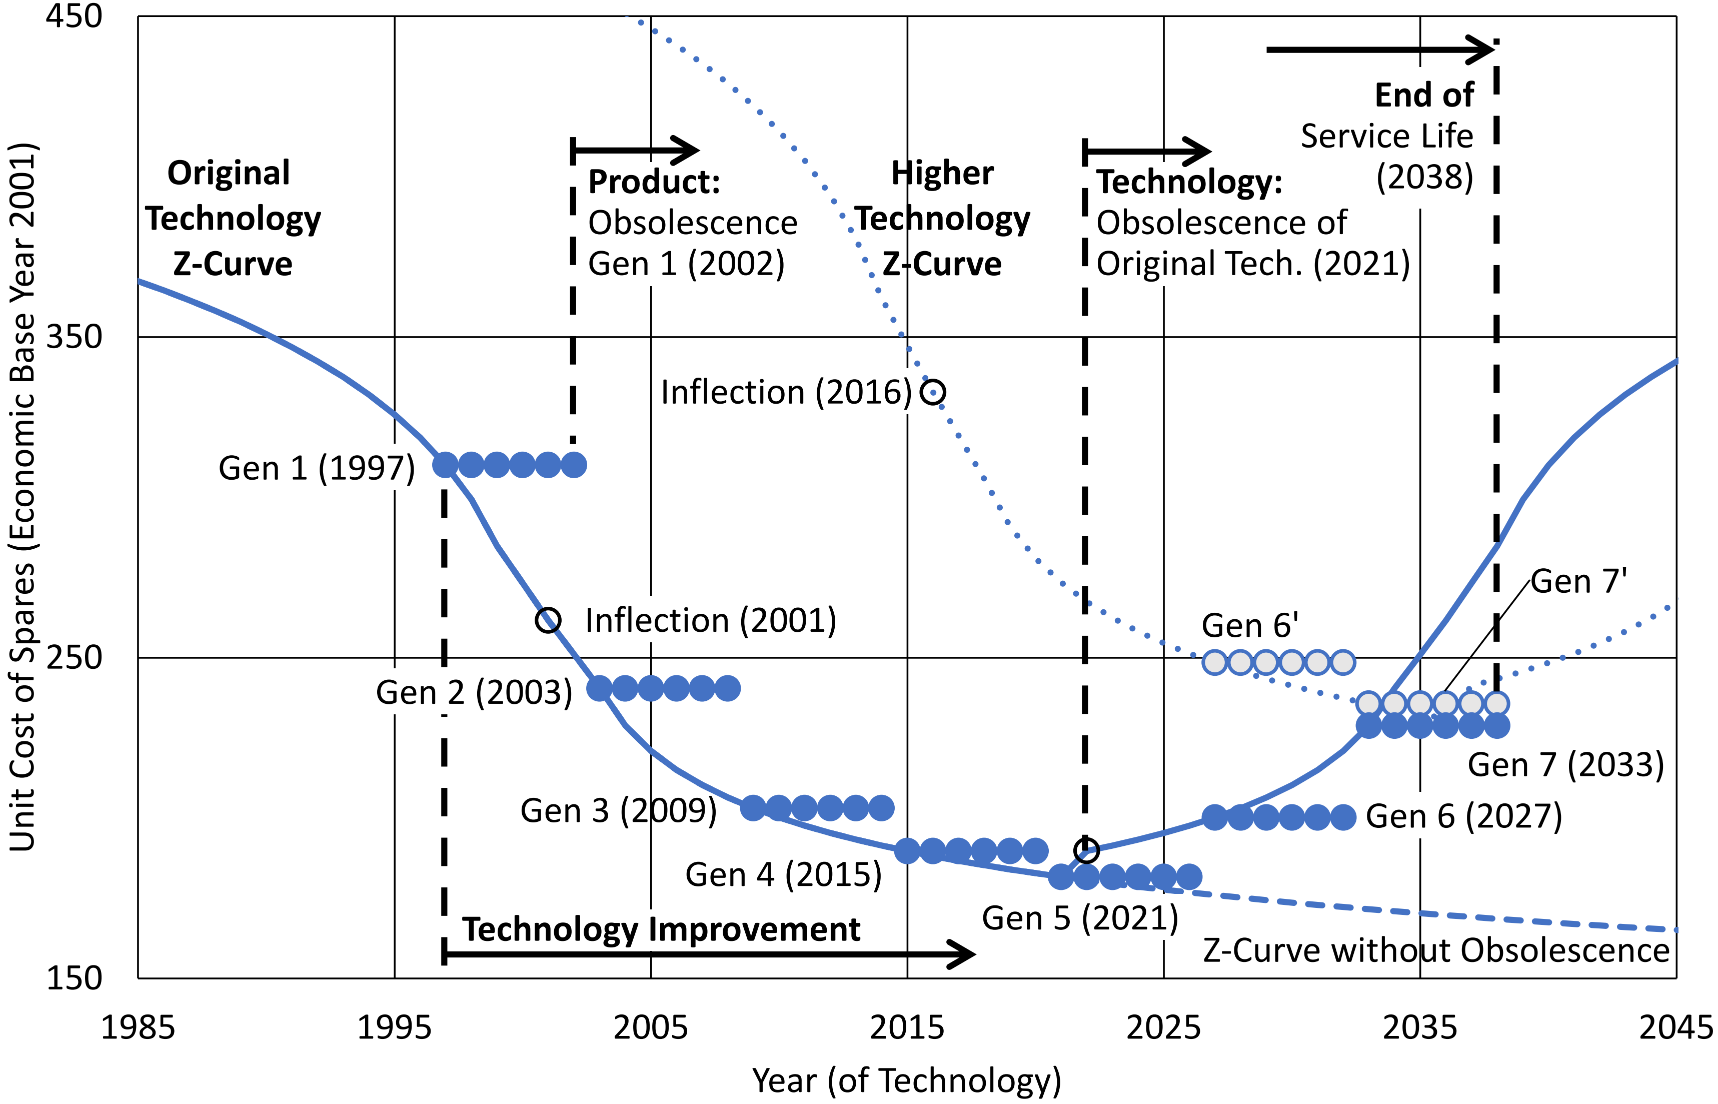

LCCM Cost Model:

- Besides the control of technology improvement, an LCCM cost model also allows **technology growth of spares**, meaning the change to a higher technology's Z-curve
- Gen 6 (2027) and Gen 7 (2033) suffer from increasing unit costs due to the original technology's obsolescence
- However, the rising curve aligns with a higher technology Z-curve (inflection 2016)
- Here is an option to replace Gen 6 and 7 with generations Gen 6' and 7' **based on a higher technology that is also 15 years younger!**

Flat Earth Fiction and LCCM Fact can be combined into one common chart, showing lots of possibilities for obsolescence mitigation

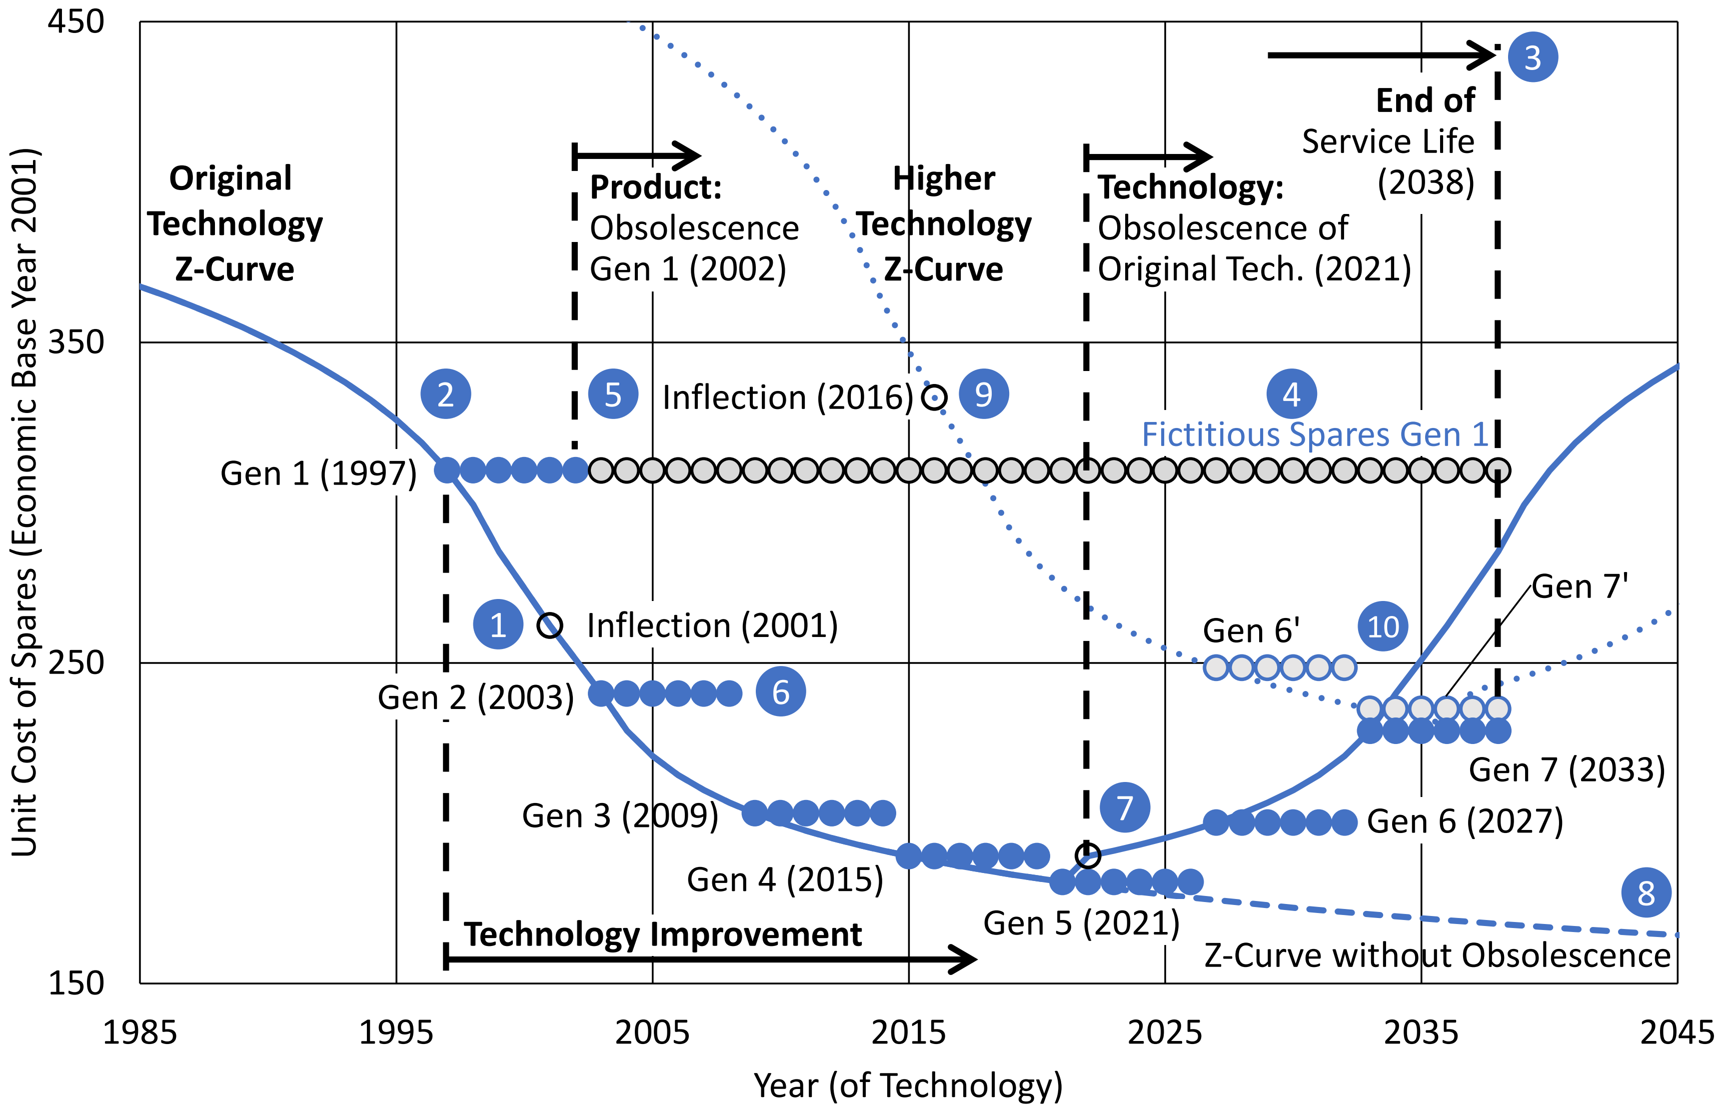

Traditional Cost Model: A product is based on a technology curve (Z-curve) with the inflection point in 2001❶. Its generation Gen 1 starts production with technology frozen in 1997❷. Service life ends in 2038❸. Spares from Gen 1 will be available for the whole 42 years❹!

LCCM Cost Model: Gen 1 **product obsolescence** occurs in 2002❺. Gen 2 replaces Gen 1 in 2003❻. Gen 3 to Gen 5 follow until 2021❼. After 2021, the **original technology becomes obsolete**, the curve deviates upwards, departing from the Z-curve minus obsolescence❽. The rising curve aligns with a higher technology Z-curve (inflection 2016)❾. Here is an option to replace Gen 6 and 7 with Gen 6' and 7' based on technology developed 15 years later❿.

A brief overview of different mitigation strategies

None (Default): The cost model will assume that all spares will remain the same for the whole service life of the system. This option is only advisable for short service lives and purely structural items.

Last Time Buy (LTB): This strategy picks a year in which all required spares will be purchased at once to support the system for its remaining service life. When stocking spare parts, permissible storage times may be limited. Furthermore, the probability of parts failing may increase due to ageing. LTB can be used as closure for the other mitigation strategies.

Equivalent in Form, Fit and Function (FFF): This strategy will replace the original product with a successor on the **same technology curve,** but more recent (**later Year of Technology**). The successor product performs the same, uses the same technology, but with the benefit of additional years **of technology improvement leading to lower unit cost.**

Technology Refresh (TechRef): This strategy replaces the original, now obsolete product with a successor on a **higher technology curve** with a **later Year of Technology**. Pursuing a higher technology may come from the need for performance improvements, increased competition, or grown customer expectations.

Product is discontinued	Substitute with same technology available?	Substitute with higher technology available?	Applicable obsolescence mitigation strategy
No	–	–	None (Default)
Yes	No	No	Last Time Buy (LTB)
Yes	Yes	No	Equivalent in Form, Fit, and Function (FFF)
Yes	No	Yes	Technology Refresh (TechRef)

Relative unit costs of spares for obsolescence mitigation with Last Time Buy (LTB, example)

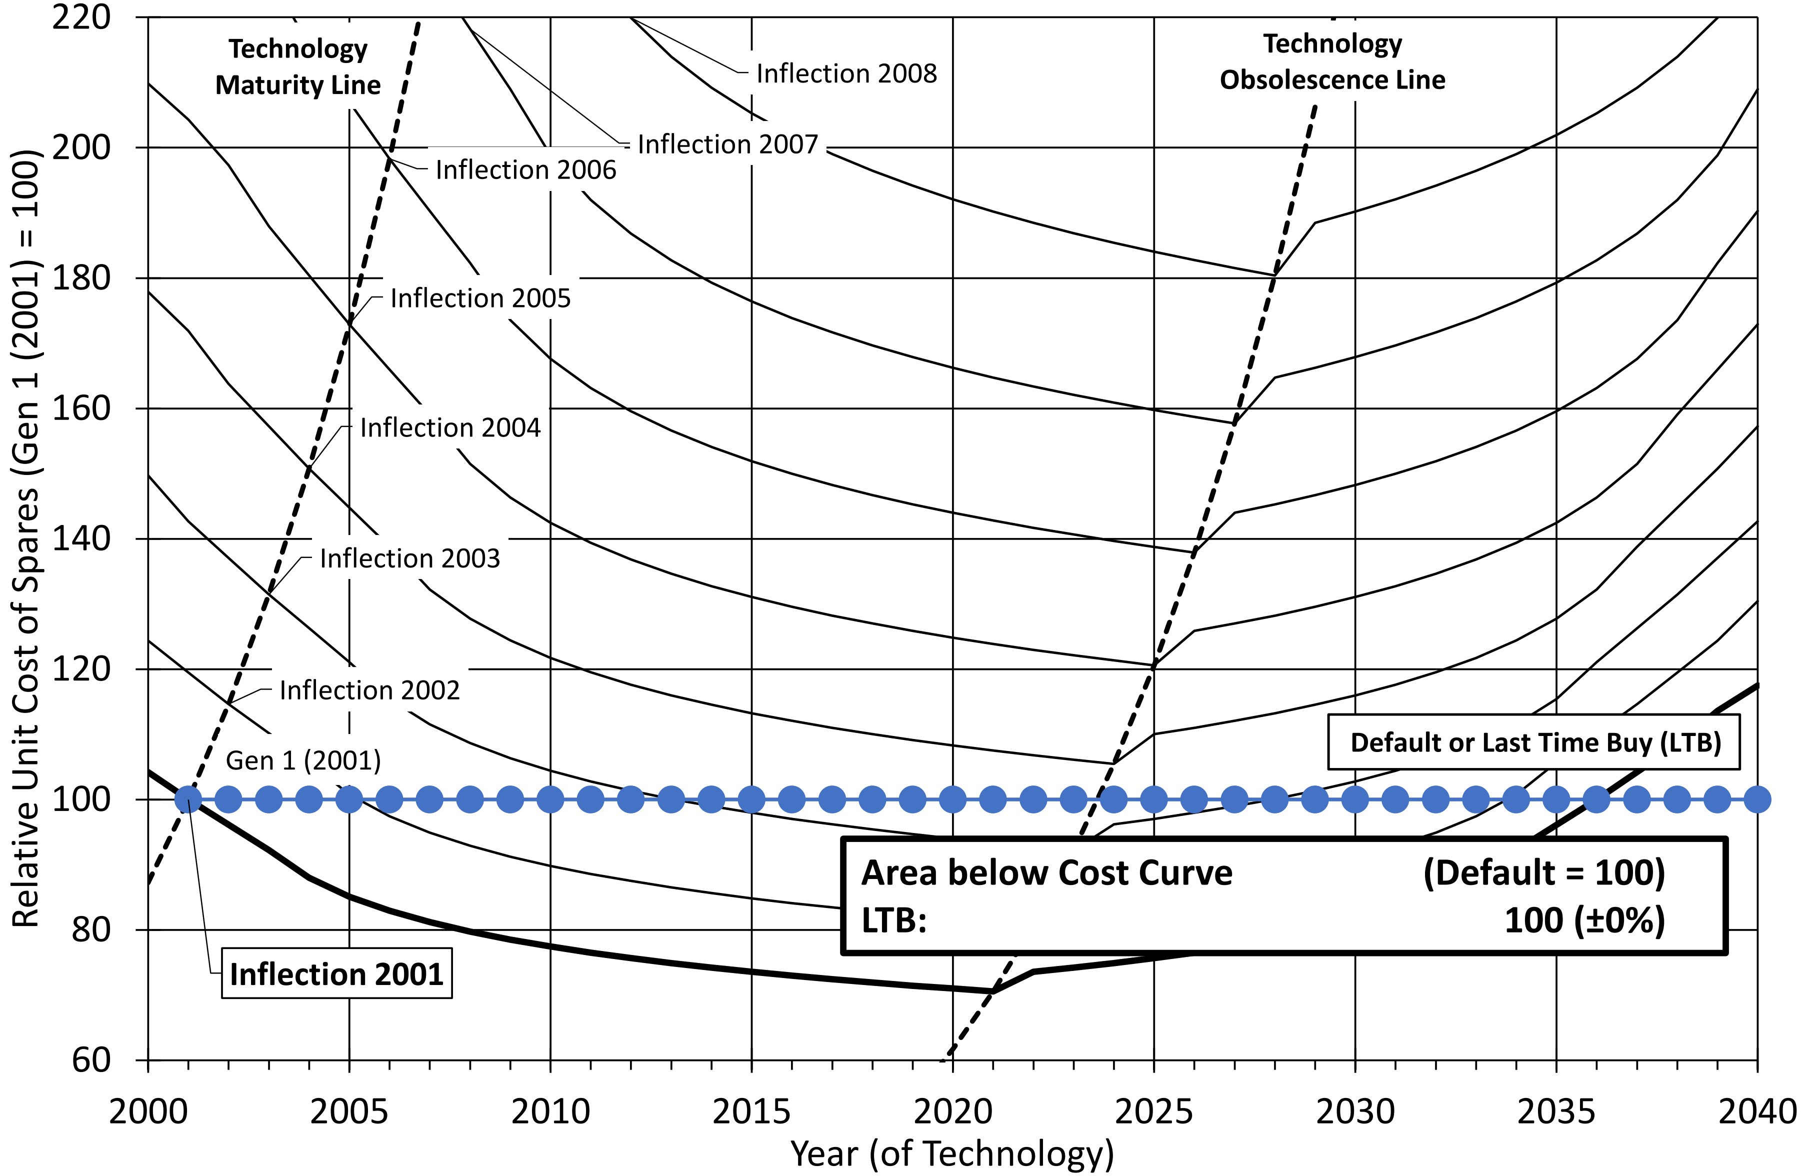

Area below Cost Curve **(Default = 100)**
LTB: **100 (±0%)**

Last Time Buy (LTB) picks a date when all spares required for the remaining service life will be purchased at once.

All replenishment spares are the same Gen 1, with the same technology curve (inflection 2001) and same year of technology. Unit cost of spares stays the same.

Latest LTB date is in the year the original product is discontinued. The outstanding spare parts requirement at the time of product obsolescence is summarised as a final stock and brought forward. The LTB stockpile will need more storage space than the default replenishment spares supply; storage space will decrease over time as LTB stock is depleted.

Relative unit costs of spares for obsolescence mitigation with Equivalent in Form, Fit, and Function (FFF, example)

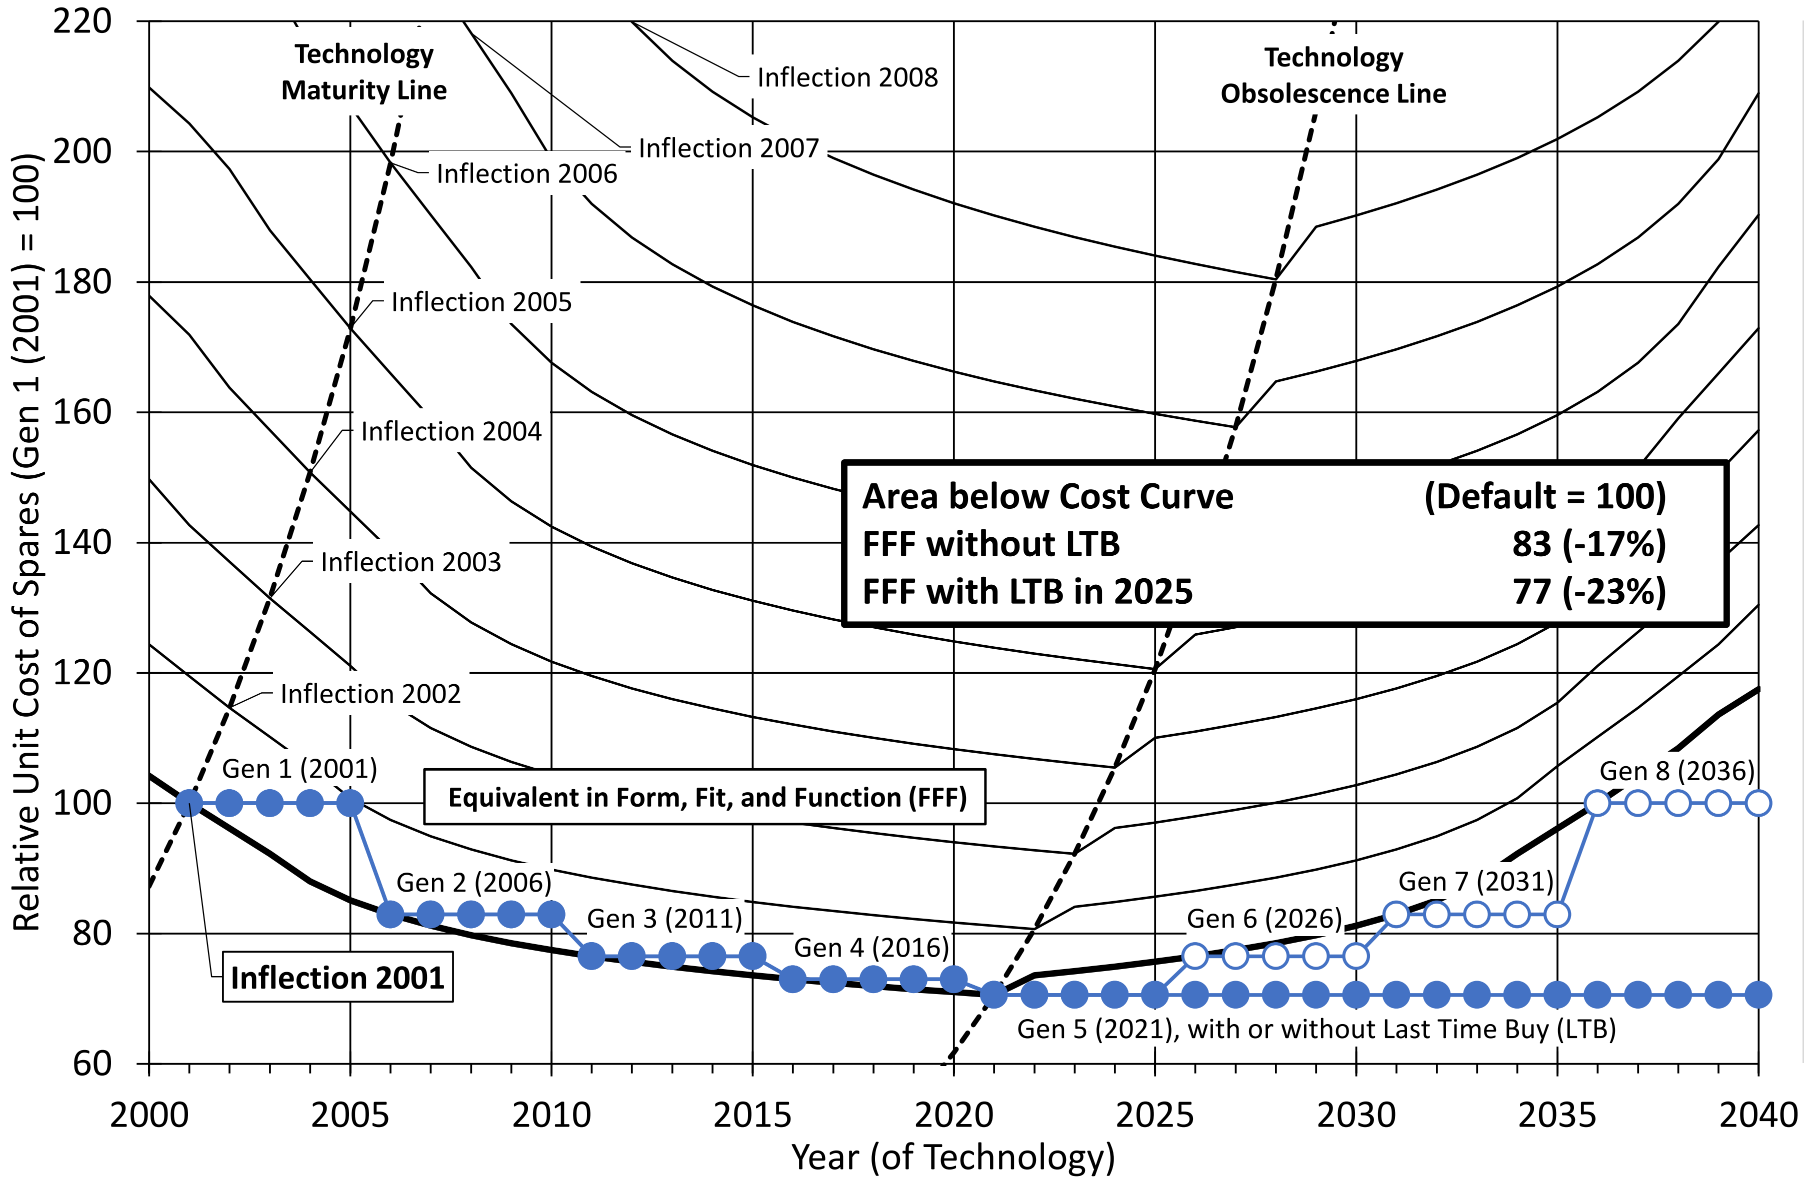

Area below Cost Curve	**(Default = 100)**
FFF without LTB	**83 (-17%)**
FFF with LTB in 2025	**77 (-23%)**

Equivalent in Form, Fit and Function (FFF) will replace the original product (Gen 1, 2001) with a more recent (later year of technology) successor product on the same technology curve (inflection 2001).

Replacements occur every five years, when a new product generation (Gen 1+x) is released.

Spares enjoy cost reduction (without inflation) due to technology improvement until technology obsolescence sets in after 2021.

When Gen 5 is discontinued in 2025, a closure with LTB would offer the lowest cost of any mitigation strategy (-23% compared to -17% without LTB).

Relative unit costs of spares for obsolescence mitigation with Technology Refresh (TechRef, example)

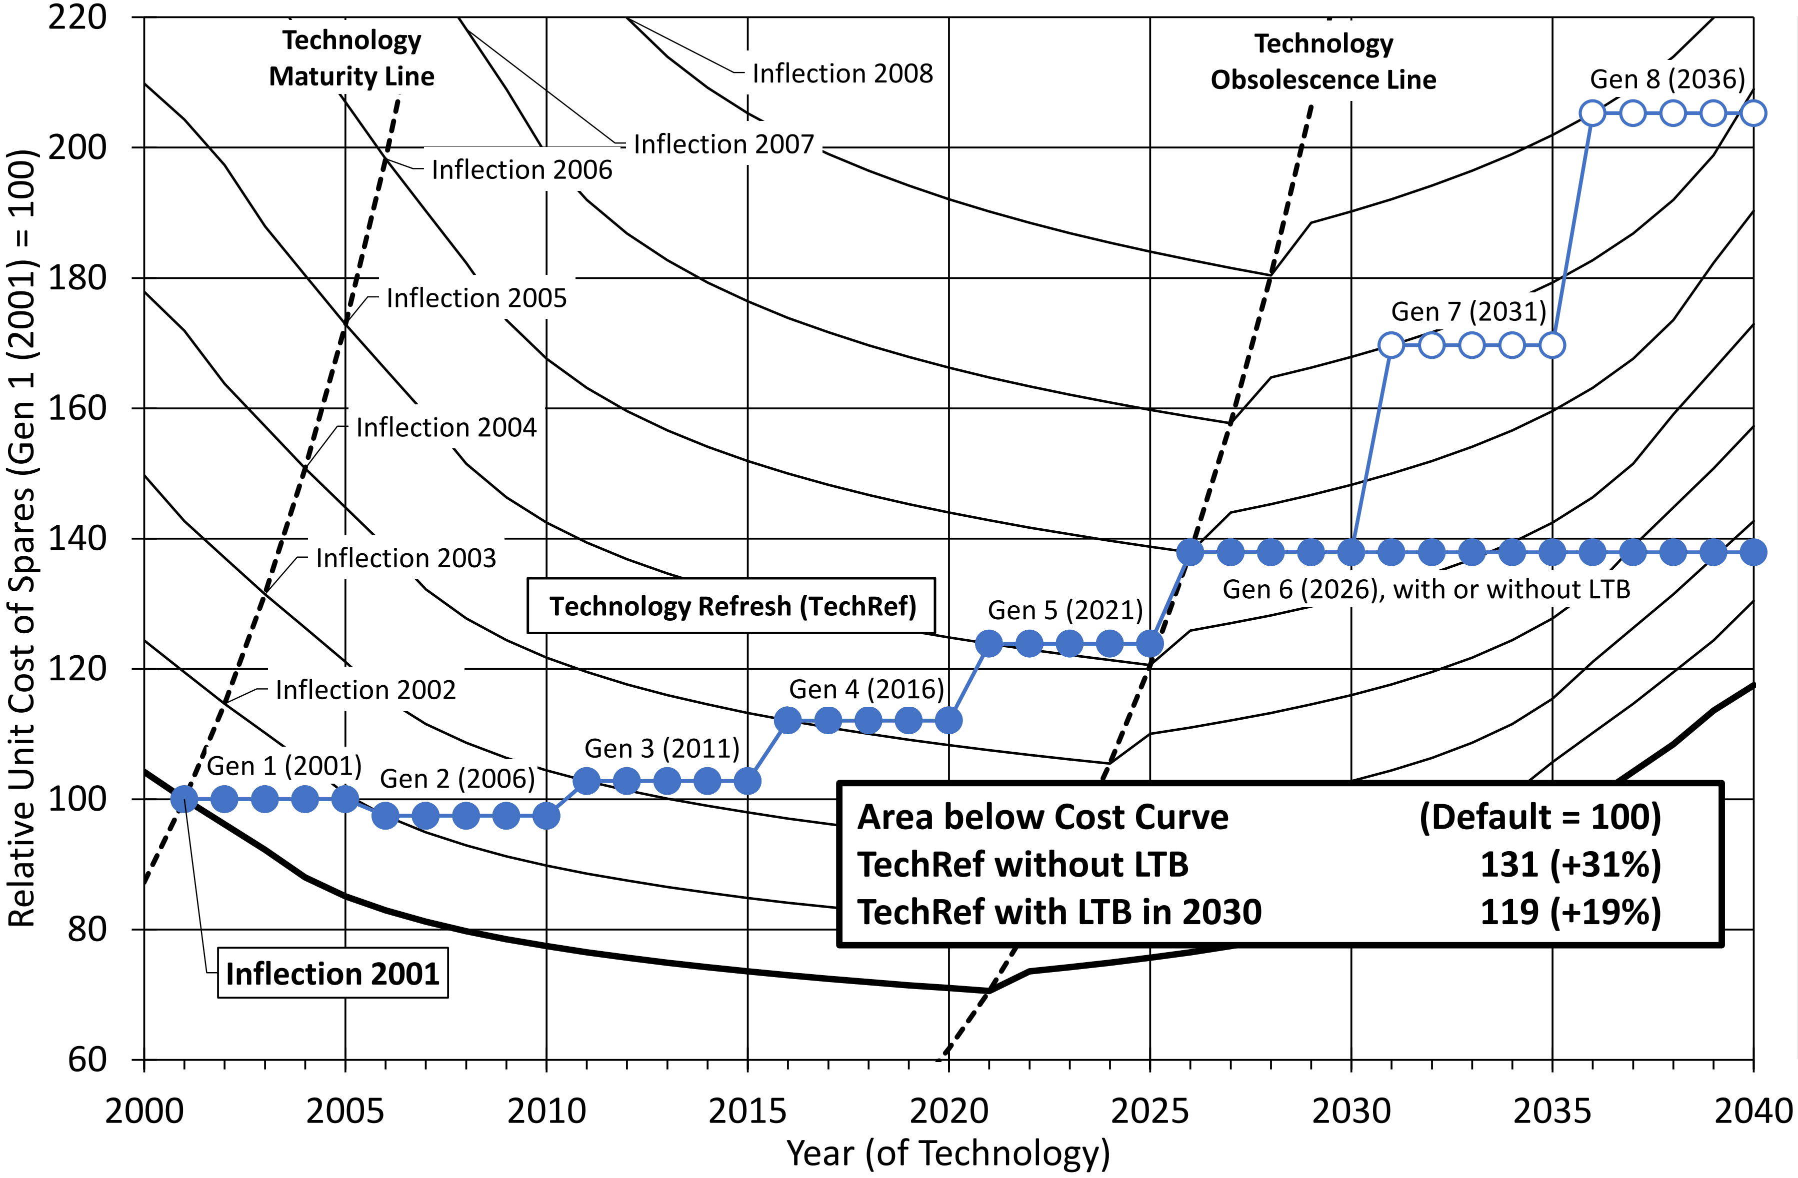

Area below Cost Curve	(Default = 100)
TechRef without LTB	131 (+31%)
TechRef with LTB in 2030	119 (+19%)

Technology Refresh (TechRef) applies technology growth as opposed to technology improvement.

TechRef has manufacturers replacing the original, now obsolete, product with a successor on a higher technology curve (inflection is later) with a later year of technology.

A refresh occurs every five years, whenever a new product generation (Gen 1+x) is released.

When Gen 6 is discontinued in 2030, a closure with LTB would offer a lower cost compared to having another two generations (+19% with LTB, compared to +31% with Gen 7 and Gen 8).

Additional input parameters for modelling various obsolescence mitigation strategies for spare parts

Cost model, traditional	Cost model, LCCM
- Technology year [calendar year] - Factor for technology improvement - Factor for technology obsolescence → **Technology obsolescence** can be modelled, but technology is frozen only once; **it is the same for production units and spares!** → **Product obsolescence** is not modelled	- Desired mitigation [LTB, FFF, TechRef] - Start of product obsolescence [calendar year] - Product cycle [years] - Technology growth with technology refresh [%] - Planned number of technology refreshments - Optional closure with Last Time Buy (yes/no) - Time of Last Time Buy [year] - Storage space for final stockpile [m^3]

Desired Obsolescence Mitigation Strategy	Last Time Buy (LTB)	Equivalent in Form, Fit and Function (FFF)	Technology Refresh (TechRef)
Start of Obsolescence Mitigation [Year]	●	●	●
Product Cycle [Years]	–	●	●
Technology Growth per Technology Refresh [%]	–	–	●
Number of Technology Refreshments	–	–	●
Closure with LTB (for other Mitigation) [yes/no]	–	●	●
Year of Last Time Buy	–	●	●
Storage space for LTB stockpile [m^3]	●	(●)	(●)

Adding product obsolescence of spare parts is the final touch to a realistic maintenance cost curve

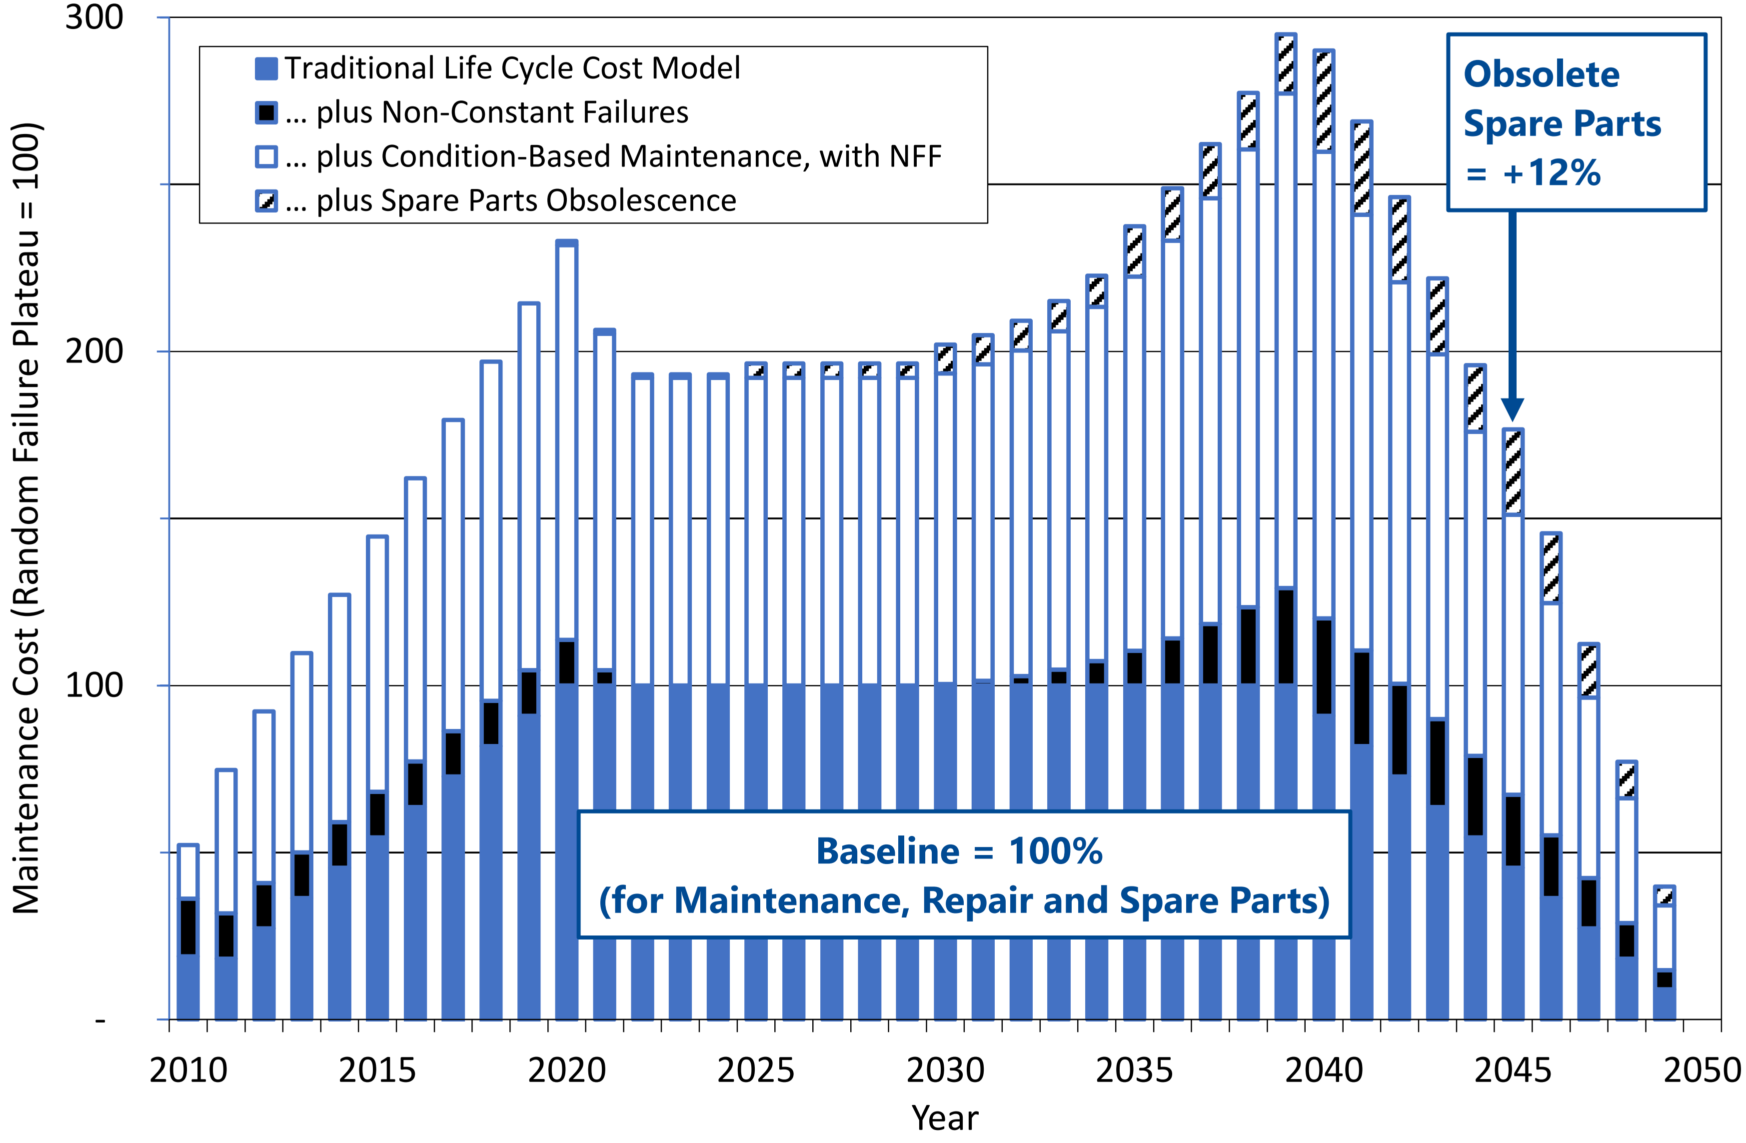

Obsolete Spare Parts = +12%

Baseline = 100%
(for Maintenance, Repair and Spare Parts)

Modern aerospace and defence systems typically boast a lengthy service life of 40 years or more, necessitating sustained support through spare parts.

The product cycles for many of these spares are notably short, with obsolescence occurring as frequently as every five years. This rapid cycle presents various mitigation options, such as last-time buys, finding equivalents in form, fit, and function, or implementing a technology refresh.

Analysing the best course of action in terms of availability and cost several years in advance is essential. Thus, a deeper understanding of the cost drivers involved in mitigating obsolescence for spare parts is vital.

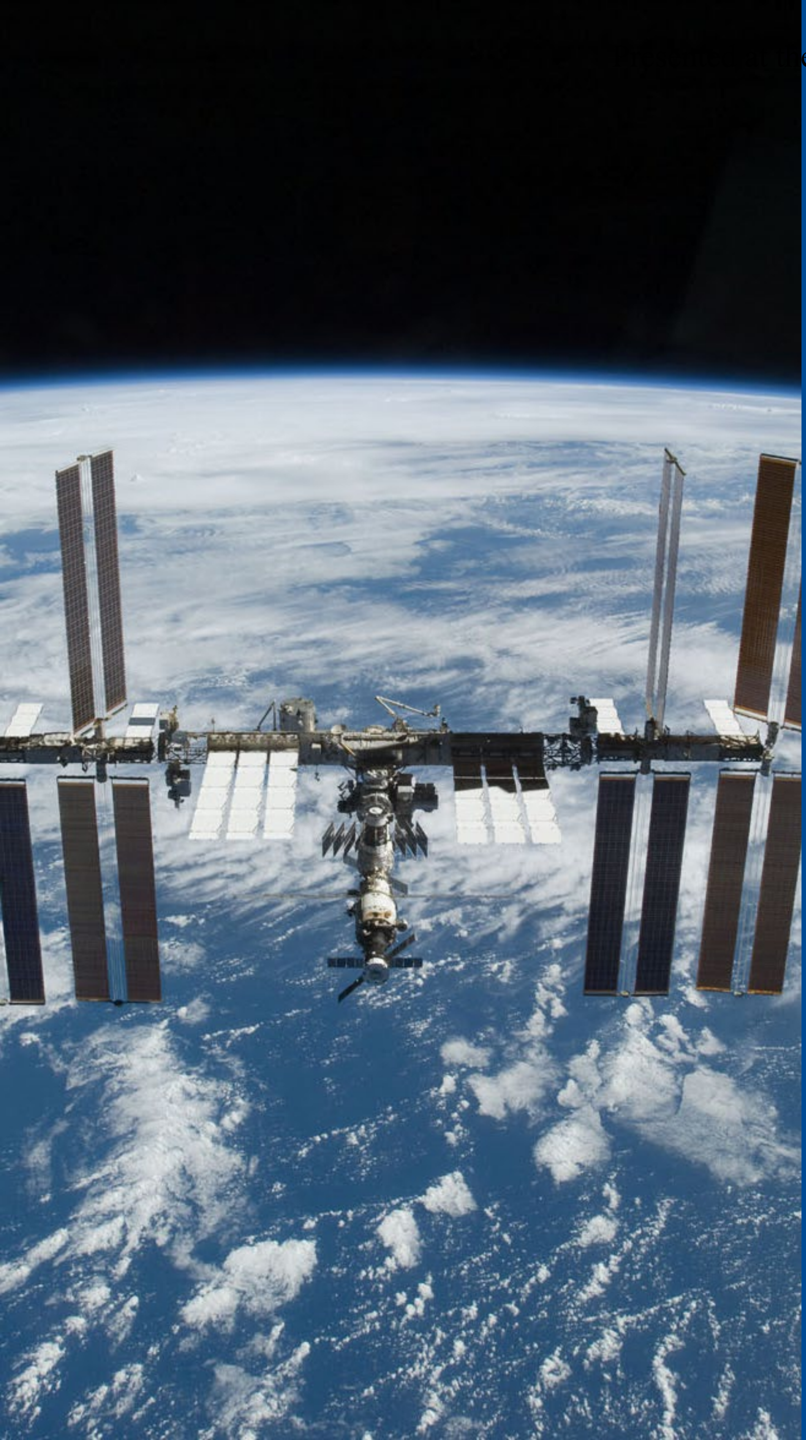

4

Earth is no longer flat.

Lessons Learned and Conclusion

LCCM-enhanced estimating leads to maintenance cost results that are significantly higher and more realistic than on “Flat Earth”

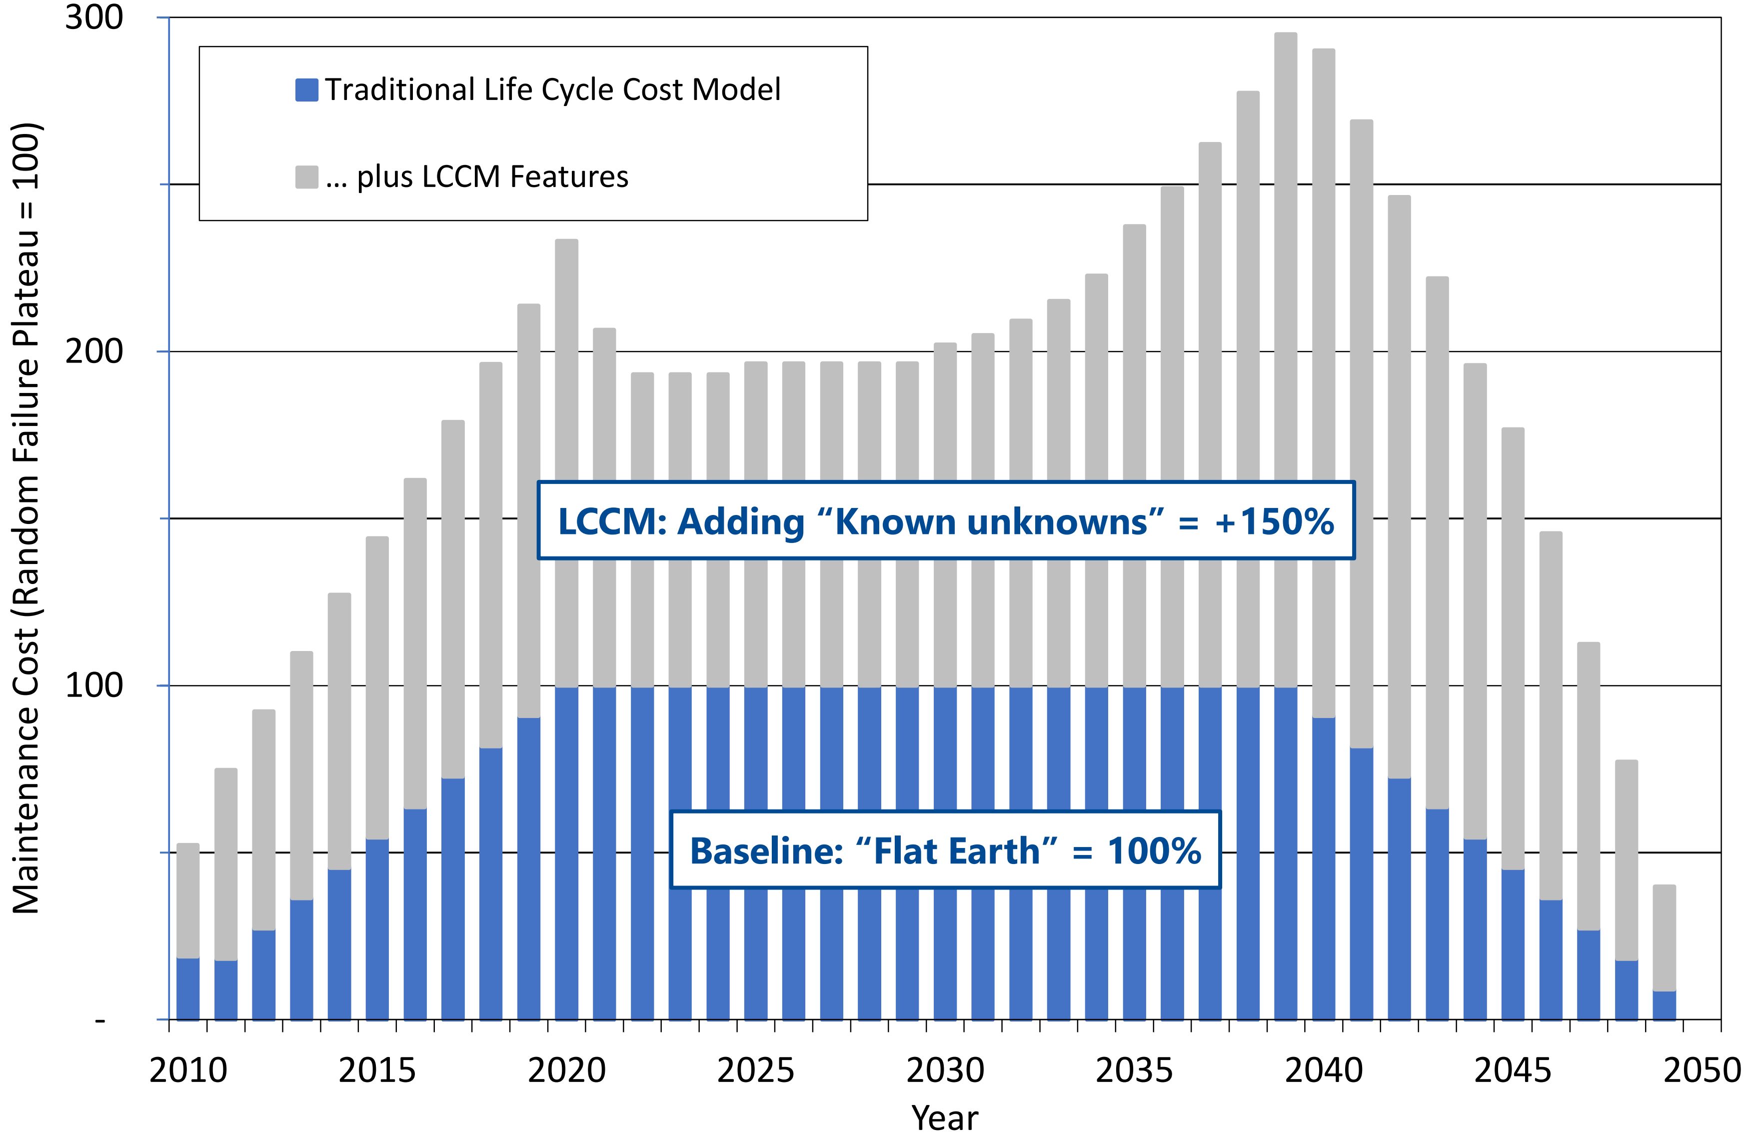

LCCM: Adding “Known unknowns” = +150%

Baseline: “Flat Earth” = 100%

In addition to traditional “Flat Earth” estimating, the following typical **increases of maintenance cost** are to be expected:

- **plus 10–70%** for additional effort due to higher failure rates coming from “infant mortality” and age effects → **more failures**
- **plus 100–300%** for additional removal, shipping, testing and reinstallation of units after “no fault found” (NFF) → **more labour effort**
- **plus 10–30%** for mitigating product obsolescence of replenishment spares → **more material cost**

Concluding Remarks

In the LCCM study presented herein, IABG delved into the top three “known unknowns” driving excessive maintenance costs and specified a model to estimate their impacts. While projecting costs 20 years or more into the future introduces a degree of uncertainty, it is crucial not to overlook these factors.

Their understanding significantly enhances the accuracy of maintenance cost estimates. Amidst data uncertainties, always remember that prioritizing **completeness over accuracy** fosters a more robust and comprehensive approach to life cycle cost (LCC) estimation.

Thank you
for asking questions

Fabian Eilingsfeld, Birgit Biehn
and Christian Schindler
Logistics & Cost Engineering DSCC13
Einsteinstrasse 20
DE-85521 Ottobrunn

E-Mail: eilingsfeld@iabg.de
biehn@iabg.de; schindlerc@iabg.de
Internet: www.iabg.de