Stiffness Logic of Bridge Design

Chris Chai

2026

Stiffness Logic of Bridge Design

Author: Chris Chai

First published on Leanpub
Version 1.0 (Spring 2026)

"*El nacimiento de una estructura es siempre una idea,
moldeadora de la materia en una forma resistente
para cumplir una misión.*"

— Eduardo Torroja, *Razón y ser de los tipos estructurales*, 1957

*(The birth of a structure is always an idea —
one that moulds matter into a resistant form to fulfil a mission.)*

Contents

Introduction

A Neglected Perspective

Every experienced bridge engineer has probably encountered this situation: the cross-section calculated to satisfy strength requirements is clearly adequate, yet the deflection check fails. The section is enlarged and rechecked — strength now has even more reserve, but deflection barely passes. The final section adopted may have a strength utilisation of only 60 per cent. The remaining 40 per cent is the cost of satisfying stiffness.

This is not an isolated phenomenon. Consider the following figures.

In a cable-stayed bridge with a main span of 600 metres, the rational girder depth based on load-carrying capacity alone is approximately 2.8 metres. To satisfy mid-span deflection limits and natural frequency requirements, the adopted depth is 3.8 metres — an increase in steel consumption of roughly 30 to 40 per cent. In a suspension bridge at a major strait crossing, the stiffening girder could be 3.0 metres deep based on strength alone; 4.2 metres is required to satisfy wind-induced stability criteria and deflection limits. A 32-metre simply-supported high-speed rail box girder requires a depth of approximately 1.8 metres for load-carrying capacity, but 2.2 metres to satisfy natural frequency and dynamic deflection limits. Even in urban viaducts — moderate spans — the theoretical reinforcement ratio based on strength is only 0.8 per cent, but to control crack width and long-term deflection the actual ratio is raised to 1.4 per cent.

The pattern is consistent: in the most active areas of contemporary bridge engineering, section dimensions are governed not by strength but by stiffness-related criteria — deflection, frequency, flutter stability, crack width.

Yet our design methods still begin with strength. Size the section for strength, check stiffness, enlarge if it fails, check again — iterate. Engineers are so accustomed to this loop that few pause to ask a seemingly simple question: if stiffness is the true governing criterion, why not begin with stiffness from the outset?

This question touches something deeper: is our understanding of structural stiffness sufficient to support a mode of thinking that starts from stiffness?

This book attempts to provide exactly that mode of thinking — not through complex mathematical derivations (that is the task of another book), but by retelling the story of bridges through the lens of stiffness. When you view arch bridges, suspension bridges, cable-stayed bridges, and girder bridges through the eyes of stiffness, much familiar engineering experience suddenly acquires new explanations; seemingly unrelated phenomena reveal deep connections; and those design intuitions that "everyone knows but no one has articulated clearly" become expressible.

One Equation, Four Logics

Understanding the stiffness logic of bridges requires only one equation:

$$K_T = K_E + K_G$$

This is a standard result in structural mechanics. K_T is total stiffness — it determines how much the structure deforms under load, how high its vibration frequencies are, and whether it is stable. K_E is section stiffness — arising from the quantity and arrangement of material in the cross-section; EI and EA both belong to this category. K_G is geometric stiffness — arising from the influence of the structure's shape on the internal-force transmission path. It is not a property of the section but a product of the overall configuration.

Geometric stiffness is the least intuitive, yet most illuminating, part of this equation.

A straight beam under lateral load can resist only through the bending capacity of its section. Where the load is applied, bending moment arises; the section must be large enough to resist it. Here geometric stiffness is zero: total stiffness equals section stiffness. What material you put in is exactly what stiffness you get — no additional help from shape.

Now bend that straight beam into an arch. Same span, same load, but the load path changes fundamentally. The curved geometry redirects vertical load into axial compression along the arch axis — the section resists not by bending but by axial compression. Under axial force the entire section participates, far more efficient than bending. But compression simultaneously introduces a risk: it drives the structure towards instability. This is negative geometric stiffness. The efficiency of the arch and its vulnerability share a single origin.

Now consider a cable. A flexible rope has no bending stiffness at all, yet under tension it resists lateral deformation — the greater the tension, the stronger the resistance. This is positive geometric stiffness. The main cable of a suspension bridge has virtually no section stiffness; all its stiffness comes from the tension established by dead load. Greater dead load, tighter cable, stiffer bridge — entirely contrary to the experience of girder-bridge engineers, where dead load is a pure burden.

With this single equation, the essential distinction among four bridge types becomes immediately clear.

Girder bridge: geometric stiffness is zero. Total stiffness equals section stiffness. What material you have is what stiffness you get — no shape to help, no shape to hinder. Simple, reliable, but span-limited.

Arch bridge: geometric stiffness is negative. Shape converts bending moment into axial compression, activating the section's axial stiffness and breaking through the span limits of girder bridges. But axial compression simultaneously produces negative geometric stiffness: the greater the load, the closer to instability. Efficiency and risk share a common source.

Suspension bridge: geometric stiffness is positive, and dominant. The contribution of cable geometric stiffness to total vertical stiffness far exceeds the elastic stiffness of the stiffening girder. Dead load is the source of stiffness.

Cable-stayed bridge: geometric stiffness and section stiffness are strongly coupled, with multiple parameters independently adjustable. Each stay cable is an independently controllable elastic support; cable spacing, tower height, cable section, and force distribution can all be actively configured. This marks the leap from passively receiving stiffness to actively planning it.

From girder to arch to suspension to cable-stayed is not a gradual refinement of the same logic, but a qualitative change in the mode of stiffness generation. To understand this is to understand the inner thread of bridge structural evolution.

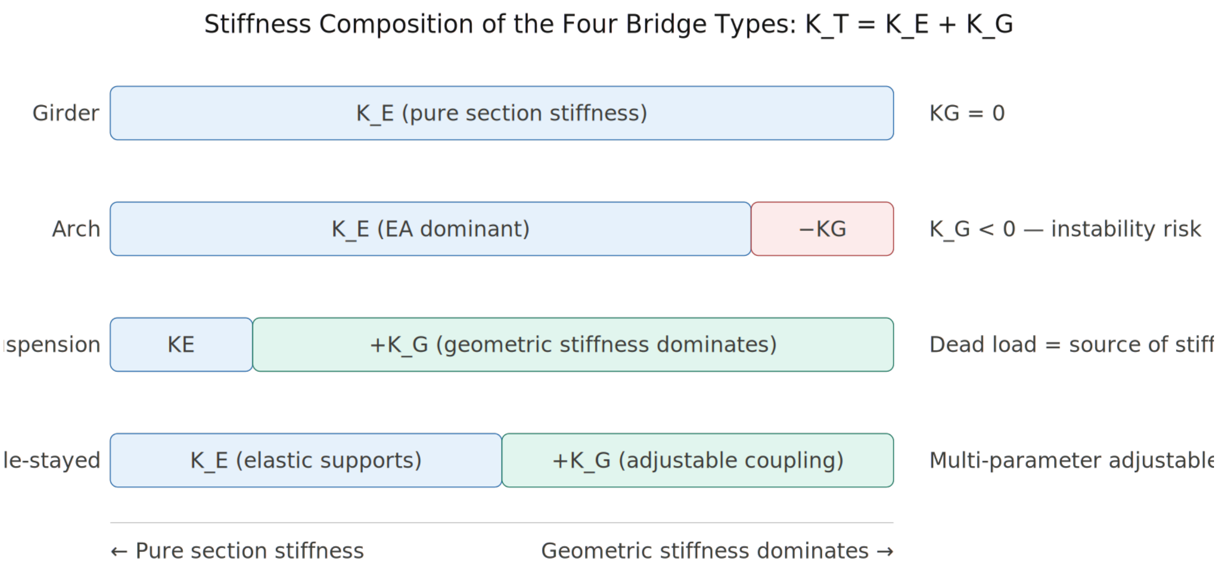

Figure 1 — Stiffness composition of the four bridge types ($K_T = K_E + K_G$). Blue bars show section stiffness K_E; green and red show positive and negative geometric stiffness K_G.

What This Book Does

The task of this book is clear: to retell the story of bridge structures through the lens of stiffness.

Part One is the core of the book. Four chapters interpret the stiffness logic of arch bridges, suspension bridges, cable-stayed bridges, and girder bridges — where stiffness comes from, which parameters govern it, and why one arrangement rather than another. Each chapter derives a few key conclusions from simplified mechanical models, enabling the reader to understand “why this bridge has adequate stiffness” or “why it does not” without finite-element analysis. These simplified models are estimation tools, not substitutes for rigorous analysis; but they provide something finite elements cannot — causal relationships among parameters.

Part Two moves beyond single systems. Composite structures — tied arches, extradosed bridges, continuous rigid frames — seek new positions between two existing stiffness logics. Boundary conditions — foundations, anchorages, bearings — profoundly affect system stiffness. Humanity has invented only four or five fundamentally different structural systems: why? The answers to all these questions find unified explanation within the stiffness framework.

Part Three widens the view. Why is the creativity of engineers systematically weakened? Is there a theoretical basis for bridge form innovation? What does the paradigm shift from strength to stiffness mean? These discussions move beyond mechanics to touch design methodology, engineering education, and broader industrial transformation.

A particular note: this book involves almost no mathematical derivation. The total number of core formulae does not exceed twenty; every one is accompanied by a complete verbal explanation, and skipping the formulae does not impair understanding. The reader needs no background in finite-element methods or advanced mechanics — a basic knowledge of bridge engineering and curiosity about structural behaviour will suffice.

Behind this book lies a complete theoretical framework. In 2025, Professor Mang of Vienna University of Technology established a rigorous variational criterion for stiffness extrema, providing for the first time a mathematical foundation for the quantitative study of stiffness. On that foundation, a series of engineering tools — coupling matrices, critical indices, energy ratios — can be developed, forming the theoretical core of what is termed “stiffness-based design”. Those contents appear in the author’s companion volume, *Stiffness-Based Design: A New Paradigm for Structural Engineering*. This book is the gateway to that framework: first see how the stiffness perspective transforms our understanding of bridges, then decide whether to explore its mathematical roots.

Before We Begin

A word on reading order.

The sequence of the four chapters in Part One is deliberately arranged: arch → suspension → cable-stayed → girder. This follows neither historical chronology nor span length, but the unfolding of stiffness logic. The arch bridge is the first invention of “shape creating stiffness”; it raises a central question — if shape can convert bending into axial compression, what about the reverse? The suspension bridge is precisely that reverse: the geometric stiffness of tension is the exact mirror image of compression. The cable-stayed bridge goes further: stiffness can not only be determined by shape but actively configured. Only after these three does the girder bridge appear — the “unassisted” state, with no geometric stiffness to help — so that the reader can fully appreciate what each of the preceding systems has broken through.

For readers with limited time, the four chapters of Part One plus Chapter 7 (“What Makes a Structure a System”) will suffice — those five chapters carry the highest density of insight in the entire book.

Let us begin with an arch bridge.

Note: All numerical examples in this book use hypothetical data for illustrative purposes only.

Part One

Four Stiffness Logics

Chapter 1

Arch Bridges: How Shape Creates Stiffness

How shape creates stiffness, and why this is a transaction with a price

1.1 A Recurring Predicament

In the design of long-span arch bridges, a troubling phenomenon recurs.

An engineer sizes the arch rib section based on material strength and load-carrying requirements, then proceeds to the deflection check and finds it fails. The section is enlarged, rechecked — it may still fail. After several iterations the section finally satisfies all requirements, but the strength check now shows enormous reserve: material load-carrying capacity is far from being fully utilised.

This is not peculiar to any individual project. From concrete-filled steel tube arches spanning 200 metres to steel-skeleton arches spanning 500 metres, from cantilever casting to cable erection, arch bridges of differing spans, materials, and construction methods all share this experience. The starting point of the design process is strength; the ultimate determinant of section size is stiffness. The mismatch between starting point and endpoint forces the engineer to go around in circles.

The problem lies not in imprecise calculation — our finite-element models are precise enough. The problem lies at a more fundamental level: we are using a bending-based way of thinking to understand a problem that is not, in its essence, a bending problem.

To understand what this means, we must first understand what arch deformation actually is.

1.2 Three Axes: The Starting Point for Understanding Arch Deformation

Understanding arch bridge deformation requires a correct geometric picture. Unlike a girder bridge, the state of an arch cannot be described by a single line; it requires three, each with a distinct physical meaning.

The first is the **stress-free fabrication line**. It is the geometry of the arch rib in the absence of any loading — including self-weight. It is the shape as manufactured in the factory: the structure's “as-shipped” form.

The second is the **design thrust line**. Under the full dead load, it is the ideal structural axis uniquely determined by equilibrium. When the arch axis coincides with the thrust line, internal forces are predominantly axial compression with minimal bending moment. For a parabolic arch under uniformly distributed dead load, this condition can be satisfied exactly.

The third is the **actual deformed profile**. It reflects the true equilibrium position under the additional effects of live load, temperature change, and construction tolerances. The deviation of this profile from the design thrust line is what the engineer must control.

With these three lines understood, two common confusions are resolved.

The precamber set during construction is not, in essence, compensating for “bending deflection”. It is raising the profile from the stress-free fabrication line up to the design thrust line under dead load. The magnitude of precamber equals the vertical distance between the design thrust line and the fabrication line.

Live-load deformation control, by contrast, keeps the deviation of the actual deformed profile relative to the design thrust line within permissible limits. The mechanism producing this deviation is not local bending deformation, but a more fundamental process: the accumulation of axial compression along the full arch.

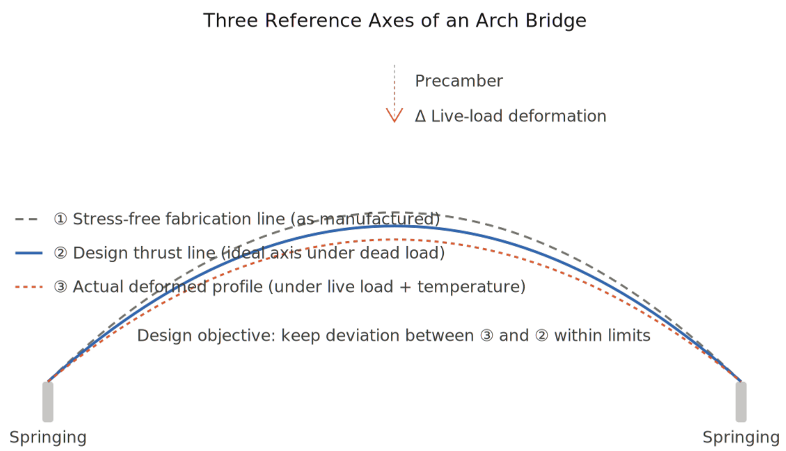

Figure 2 — Three reference axes of an arch bridge. Grey dashed: stress-free fabrication line (precambered); solid blue: design thrust line under dead load; red dot-dash: actual deformed profile under live load.

1.3 Arch Deflection Is Not Bending — A Misunderstood Problem

This is the most critical section of this chapter. If only one thing is to be retained, let it be the conclusion of this section.

The misleading bending framework

In textbooks and engineering software, an arch is typically treated as a “curved beam with initial curvature”, and deflection is understood as the result of bending deformation. Under this framework, designers naturally emphasise the role of flexural stiffness EI and treat increasing the section moment of inertia as the primary means of controlling deflection.

But this framework is misleading. For a properly proportioned arch — where the arch axis is close to the thrust line — bending moments are inherently small. What truly governs deflection is not bending deformation but the accumulation of axial shortening along the full arch. We may call this “geometric reconstitution”.

Three-step derivation: from axial force to crown displacement

Consider a parabolic arch with rise-to-span ratio f/L under uniformly distributed load q. The process can be understood in three steps.

Step 1 — Load becomes axial force. The curved geometry redirects vertical load into axial compression along the arch axis: $N \approx qL^2/(8f)$. Note the identity of form with the suspension bridge cable force $H = gL^2/(8f)$ — compression and tension, same formula, opposite signs. This is the mathematical root of the arch–cable mirror symmetry.

Step 2 — Axial force accumulates as shortening. Axial force N produces strain $\varepsilon = N/(EA)$ at every point. Individual strain is small (10^{-4} to 10^{-3}), but the full arch length S can reach hundreds of metres. Tiny strains integrated over the full length produce substantial total shortening.

Step 3 — Shortening converts to crown settlement. Through geometric compatibility — fixed springings, shorter rib — total axial shortening converts into vertical crown displacement. This process is geometric, not material: even in the complete absence of bending deformation, axial shortening alone causes the crown to settle.

Combining the three steps:

$$\Delta/L \;\approx\; C \cdot (q/E) \cdot (L/A) \cdot (L/f)^2$$

where $C \approx 0.017$ for typical engineering rise-to-span ratios (f/L = 1/4 to 1/6), varying less than ±4 per cent — effectively a constant. The core message: deflection-to-span ratio is governed by $(L/f)^2$ and L/A; C itself is not a variable that requires attention.

What this formula reveals

First, deflection-to-span ratio is inversely proportional to sectional area A and independent of moment of inertia I. The symbol I does not appear on the right-hand side. This is the most important conclusion of this chapter: **controlling arch deflection depends on EA, not EI.** Increasing the moment of inertia (e.g. deepening the section without adding area) has almost no effect on the accumulation of axial compression; increasing sectional area (directly raising EA) is what effectively reduces deflection.

Second, deflection-to-span ratio is proportional to $(L/f)^2$. The rise-to-span ratio f/L appears in the denominator squared — the most sensitive parameter controlling arch deflection. Halving the rise-to-span ratio quadruples the deflection.

A striking comparison

It is worth placing the deflection formulae of arch and girder bridges side by side.

A girder bridge: deflection proportional to L^4/EI — the fourth power of span. Double the span, deflection increases sixteenfold. This is the "L^4 law", the fundamental cause of girder bridge span limits, discussed in detail in Chapter 4.

An arch bridge: deflection proportional to L^2/EA — the second power of span. Double the span, deflection increases only fourfold.

From L^4 to L^2, the span exponent is reduced by two orders. This is not a marginal improvement within the same framework but a change in order of magnitude achieved by changing the mechanical load path. This is the secret of how arch bridges cross hundreds of metres with far smaller sections than girder bridges — the arch is not a "better beam" but a fundamentally different mode of stiffness generation. It is humanity's first breakthrough against the L^4 law.

Numbers for a typical arch bridge

Consider a typical long-span concrete-filled steel tube arch: span 400 metres, rise-to-span ratio 1/5 (rise 80 m), arch rib sectional area 4 m^2, composite elastic modulus 35 GPa, dead-load line density 300 kN/m.

Arch rib axial force: 75,000 kN. Axial strain: 5.4×10^{-4} — indeed very small. Full arch length approximately 443 m. Total axial shortening: 0.24 m — nearly a quarter of a metre. Crown vertical displacement: approximately 0.15 m. Deflection-to-span ratio approximately 1/2,700 — well within limits.

Now a thought experiment: halve the sectional area. Axial strain doubles; crown displacement becomes 0.30 m; deflection-to-span ratio drops to 1/1,350. Superimpose live load and temperature effects, and the typical control limit will very likely be breached.

Here is the critical comparison: if instead of adding area we increase the moment of inertia — say by deepening the section — the effect on this deflection is virtually nil, because the moment of inertia does not appear in the deflection formula. This is the root cause of the puzzling phenomenon in design practice: repeatedly enlarging EI within a bending-based framework, yet deflection stubbornly refuses to comply.

1.4 EA Primary, EI Secondary — Layered Stiffness Control

Based on the preceding analysis, stiffness plays a clearly tiered role in arch design.

First tier: axial stiffness EA governs deflection. It directly controls the level of axial strain and is the primary parameter determining the magnitude of geometric reconstitution. The effect of increasing EA on deflection control is linear: double the area, halve the deflection.

Second tier: flexural stiffness EI ensures stability. Under a well-fitted arch axis, EI contributes little directly to deflection. But it has a critical indirect role: under high axial compression, EI raises the Euler critical load and suppresses second-order amplification.

An intuitive picture: axial compression causes the arch axis to deviate from the design thrust line → deviation generates additional bending moment → additional moment increases deviation → a vicious cycle forms. EI is the brake that interrupts this cycle — when sufficiently large, deviation is quickly suppressed and the cycle cannot start.

In one sentence: EA exploits the advantage of geometry; EI controls the risk that geometry introduces. Neither is dispensable.

Third tier: strength is generally satisfied automatically. The sectional area required to control deflection — a global constraint governed by strain accumulation over the full arch — far exceeds what strength alone requires, which is merely a local constraint. Global constraints are always more stringent than local ones. This explains why arch bridge strength checks almost invariably show ample reserve.

This hierarchy dictates the rational design sequence for arch structures: first ensure EA satisfies deflection requirements; then ensure EI satisfies stability requirements; finally check strength. This is the exact reverse of the traditional “size for strength, then check stiffness” — and it is the concrete embodiment of stiffness-led design thinking applied to arch bridges.

The most direct practical implication: when an arch bridge fails its deflection check, do not instinctively reach for more EI. First check whether EA is adequate. Starting from EA is treating the cause; enlarging EI within a bending framework is going around in circles.

1.5 Rise-to-Span Ratio: The Result of Multiple Competing Constraints

The squared control of deflection by the rise-to-span ratio f/L is only part of its influence. It simultaneously governs several of the most important indicators of an arch bridge, and in conflicting directions — which is why engineering practice concentrates rise-to-span ratios in the narrow range of 1/5 to 1/6, rather than “the larger the better” or “the smaller the better”.

Criterion	Effect of ↑ f/L	Mechanism	Preferred direction

Crown deflection	Decreases (favourable)	N decreases, ΔS decreases	Larger is better
Springing thrust	Decreases (favourable)	$H = qL^2/(8f)$; larger f → smaller H	Larger is better
Axial compression	Decreases (favourable)	N proportional to H	Larger is better
In-plane stability	Decreases (unfavourable)	Longer rib, lower critical load	Must not be too large
Construction difficulty	Increases (unfavourable)	Greater arch height, harder erection	Must not be too large
Clearance height	Increases (unfavourable)	Crown headroom limits	Must not be too large

The first three rows push towards larger rise-to-span ratios; the last three impose constraints. The range 1/5 to 1/6 is the intersection of three competing requirements: EA gain, thrust penalty, and construction feasibility.

Comparing three rise-to-span ratios for a 400-metre arch bridge:

Criterion	f/L = 1/4	f/L = 1/5 (baseline)	f/L = 1/6
Rise (m)	100	80	67
Springing thrust ratio	0.63	1.00	1.50
Crown deflection ratio	0.64	1.00	2.25
Full arch length ratio	0.96	1.00	1.05

From 1/5 to 1/6: crown deflection increases by 125 per cent and springing thrust by 50 per cent. From 1/5 to 1/4: deflection falls by 36 per cent and thrust also decreases, but in-plane stability and construction difficulty increase markedly. The ratio 1/5 is the balance point.

The above are trend estimates from simplified formulae. The truly compelling evidence is the collective choice of engineering practice:

Bridge	Main span (m)	Actual f/L	Completion
Lupu Bridge, Shanghai	550	≈ 1/5.3	2003
Chaotianmen Bridge, Chongqing	552	≈ 1/5.0	2009
New River Gorge Bridge, USA	518	≈ 1/5.5	1977
Sydney Harbour Bridge	503	≈ 1/5.6	1932
Wanxian Yangtze River Bridge	420	≈ 1/5.0	1997

Spanning nearly a century, from different countries and design teams, these long-span arch bridges cluster without exception in the narrow band between 1/5 and 1/6. This is not the outcome of code requirements but the collective empirical confirmation of the multiple trade-offs predicted by theory.

1.6 Confirmation from Engineering Evolution

If the preceding analysis is correct — the key to controlling arch deflection lies in EA — then the successful long-span arch bridges in engineering practice should be those structural forms with ample reserves of axial stiffness. And so they are.

The **concrete-filled steel tube arch** was originally motivated by construction needs: using the steel tube as a self-supporting system during the arching phase, completing closure without scaffolding, and infilling with concrete to increase load-carrying capacity. But from a stiffness perspective, its key feature is that the compressive sectional area is significantly larger than that of a conventional box-section rib: the steel tube and core concrete work together, yielding extremely high axial stiffness EA in the completed state. The direct effect of this high EA is not primarily an increase in load-carrying capacity, but the suppression of axial compression deformation — reduced strain under live load, limited thrust-line deviation, and restrained additional bending moments.

The **steel-skeleton arch** likewise originates from construction needs — the steel truss skeleton forms a stable spatial structure before concrete is cast. After completion, the skeleton and concrete work together, creating a composite system with larger sectional area and higher axial stiffness. The Tian’e Longtan Bridge (main span 600 m) and Pingnan Third Bridge (main span 575 m), both completed in recent years, are representatives of this approach.

The common thread emerges: both forms, through their choice of structural configuration, significantly raised the axial stiffness of the arch rib without explicitly invoking “stiffness design”. This is a path of “practice first, mechanics later” — engineers found the right direction through experience and intuition; this chapter merely lays out the underlying mechanical logic systematically.

One further point: the stiffness of a long-span arch bridge does not appear suddenly at the moment of completion; it is generated progressively through the construction process. In a steel-skeleton arch, system stiffness at skeleton closure comes from the steel EA; as concrete is cast in stages, sectional area grows, EA rises, while self-weight also increases and axial force grows accordingly. Every step during construction is a race between EA gain and axial force growth. Each pour of concrete permanently alters the stiffness distribution — the construction path has a “lock-in” effect on stiffness.

1.7 Arch and Cable: Two Faces of a Mirror

From a wider perspective, the arch and the suspension cable possess a deep symmetry. Both are axially-dominated structures whose principal deformation mechanism arises from axial strain rather than bending. Increasing EA is the shared core strategy for controlling geometric deformation. Yet in geometric stiffness they display precisely opposite characteristics.

Feature	Arch Bridge	Suspension Bridge
Axial force	Compression (N)	Tension (H)
Core formula	$N = qL^2/(8f)$	$H = gL^2/(8f)$
Geometric stiffness	$K_G < 0$ (detrimental)	$K_G > 0$ (beneficial)
Role of dead load	Enemy (increases compression)	Source of stiffness
Key form parameter	Rise-to-span ratio f/L	Sag-to-span ratio α
Primary stiffness param.	EA (axial compression)	EA (axial tension)
Role of EI	Compensates negative K_G	Enabling: distribution + torsion
Strength status	Generally auto-satisfied	Generally auto-satisfied
Self-stabilising	No (greater load → danger)	Yes (greater load → stiffer)

This table is a complete mirror image of the two systems. The core formulae share the same form, but axial force has opposite signs, from which derive the positive–negative opposition of geometric stiffness, the ally–enemy distinction in the role of dead load, and the presence or absence of self-stabilisation. Everything that must be “guarded against” in arch design — negative geometric stiffness, instability risk, the detrimental effect of dead load — becomes something to be “exploited” in suspension bridge design, and vice versa.

The next chapter enters the other side of this mirror. If the dead load of an arch bridge is the enemy — the heavier, the more dangerous — then the suspension bridge will present an entirely opposite picture: dead load as the source of stiffness, the stiffening girder not for carrying load but for “enabling”, deflection even independent of section size. These counter-intuitive conclusions were answered more than 80 years ago by Li Guohao’s equivalent beam model.

1.8 What the Arch Teaches Us

The stiffness logic of arch bridges can be summarised in three statements.

First, arch deflection is fundamentally a problem of geometric reconstitution caused by accumulated axial shortening along the full arch — not a bending problem. For deflection control, EA is the primary parameter; EI principally ensures stability rather than directly controlling deflection.

Second, the rise-to-span ratio is the core form parameter of an arch bridge, simultaneously governing axial force magnitude, deflection order of magnitude, springing thrust, and stability risk. Its concentration between 1/5 and 1/6 is the result of multiple constraints acting together, not an arbitrary empirical rule.

Third, the arch bridge is the first invention of “shape creating stiffness”, but this creation comes at a price: axial compression is both the source of efficient load transfer (EA replacing EI) and the agent that systematically erodes total stiffness ($K_G < 0$). The advantages and limitations of the arch share a single origin — the first engineering instance of the “bundled transaction”.

But the arch bridge is more than a structural type. It is a teaching case about how to correctly identify the governing design factor. When we understand a problem through the wrong framework — treating deflection that is fundamentally governed by EA as a bending problem governed by EI — even the most precise calculations merely go around in circles. Finite elements will faithfully give you numbers, but they will not tell you what those numbers principally depend on. Identifying the true governing factor may be more important than computing it precisely.

The story of the arch also reveals a deeper law: the advantages and limitations of a structure share a single origin. The efficiency of the arch — replacing bending with axial compression, achieving a reduction in span exponent — and its fragility — the instability risk from negative geometric stiffness — both arise from the same thing: axial compression. Accept the efficiency that geometry brings, and you must accept the risk it brings. This is the “bundled transaction” that Chapter 7 will explore in depth — every structural system is an inseparable unity of advantages and limitations.

Chapter 2

Suspension Bridges: Dead Load as the Source of Stiffness

From Li Guohao's second-order theory to stiffness-led design thinking

2.1 Where Does the Stiffness of a Suspension Bridge Come From?

The stiffening girder of a suspension bridge is far shallower than a continuous girder of the same span. The Xihoumen Bridge has a main span of 1,650 metres, yet its steel box girder is only 3.5 metres deep — a depth-to-span ratio of approximately 1/470. A continuous girder of the same span would require at least 1/50 to 1/80. This means the flexural stiffness EI of the stiffening girder is only a few hundredths of one per cent of that of a comparable continuous girder.

Yet the deflection-to-span ratio of a suspension bridge is comparable to that of a continuous girder bridge — both in the range 1/300 to 1/500.

A bridge whose flexural stiffness is only a fraction of its counterpart, yet achieving comparable deflection performance — this enormous disparity leads to a core question: where does the stiffness of a suspension bridge come from?

The answer, of course, is the cable. But a cable is a flexible member with no bending stiffness whatsoever — lay a steel wire rope across two supports and it simply sags; it has no capacity to resist bending. The cable provides stiffness in an entirely different way: through the constraint that tension imposes on deformation. Once the cable is tightened by dead load, any vertical deformation causes the cable profile to deviate from equilibrium, and tension immediately produces a restoring force — pulling the deformation back. This stiffness, arising from the coupling of internal force and deformed geometry, is called geometric stiffness in structural mechanics. Its essence is the "stiffness-enabling" effect of the gravitational field on flexible matter.

If the arch bridge was the first invention of "shape creating stiffness" — converting bending into axial compression through curvature, achieving the span-exponent reduction from L^4/EI to L^2/EA — then the suspension bridge goes further still: it almost entirely abandons section stiffness and entrusts the fate of the structure to geometric stiffness. This is the precise mirror image of the arch, and the most extreme position on the stiffness spectrum.

2.2 The Equivalent Beam: Li Guohao's Brilliant Insight

In 1940, Li Guohao completed his doctoral dissertation at Darmstadt University of Technology in Germany, in which he made a remarkably elegant observation: the bending differential equation of the stiffening girder shows that its structural behaviour is equivalent to a beam carrying vertical load while simultaneously subjected to the cable horizontal force H as axial tension. This “tension-bending beam” is the equivalent beam.

A point that is easily misunderstood deserves clarification: the axial tension H in the equivalent beam is precisely the mathematical expression of the main cable’s action. The cable has not disappeared from the model; it is implicitly present in the form of its horizontal component H, acting as axial tension in the boundary conditions of the equivalent beam. The equivalent beam does not neglect the cable — it expresses the cable’s effect in the most concise possible way.

Figure 3 — Equivalent Beam Principle of a Suspension Bridge

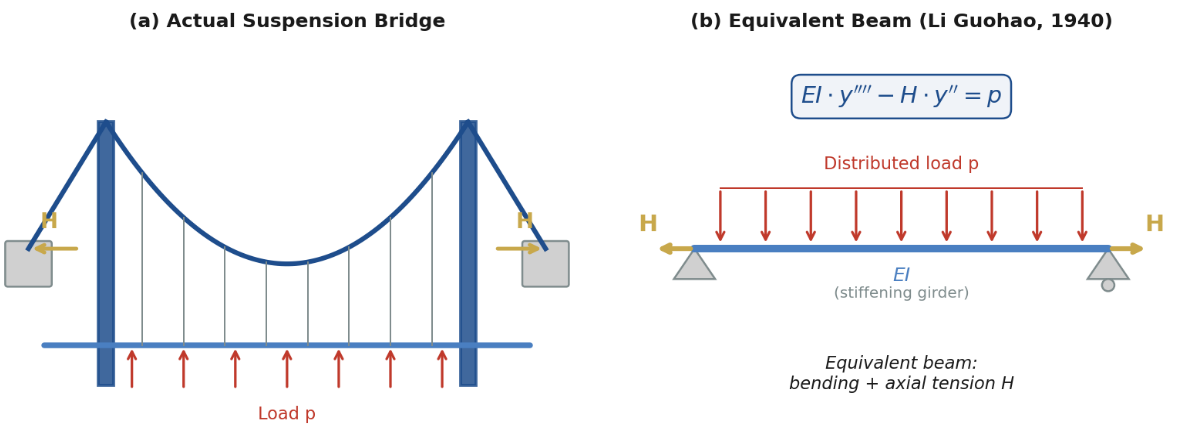

Figure 3 — Equivalent beam principle of a suspension bridge (Li Guohao, 1940). Left: actual suspension bridge with main cable, hangers, and horizontal force H. Right: the equivalent tension-bending beam, where cable horizontal force H appears as axial tension in the governing equation $EI \cdot y''' - H \cdot y'' = p$.

This equivalence means the stiffness of a suspension bridge has two distinct sources: first, the flexural stiffness EI of the stiffening girder itself — this is elastic stiffness; second, the restraining effect produced by the cable horizontal force H through the geometry of deformation — this is geometric stiffness.

Cable horizontal force and sag-to-span ratio

Under dead load g the cable forms a parabola, and its horizontal component H can be expressed by a simple formula:

$$H = gL^2 / (8f) = gL / (8\alpha)$$

where L is the main span, f is the cable sag, and $\alpha = f/L$ is the sag-to-span ratio. The meaning of this formula bears repeated emphasis: cable horizontal force is proportional to dead load and inversely proportional to sag-to-span ratio.

Since H is the primary provider of stiffness, this means: the stiffness of a suspension bridge derives fundamentally from dead load. Greater dead load and smaller sag-to-span ratio produce greater H and higher stiffness. This is entirely contrary to the experience of girder bridge engineers, for whom dead load is a pure burden.

Orders of magnitude compared

To build intuition, consider a suspension bridge with main span L = 1,000 m, sag-to-span ratio α = 1/10, dead load g = 200 kN/m. Cable horizontal force H = 250,000 kN.

If only the stiffening girder's own stiffness were available (a steel box girder roughly 3 m deep and 30 m wide, with moment of inertia approximately 2 m^4 and EI approximately 4×10^8 $kN{\cdot}m^2$), the midspan deflection under a 1,000 kN concentrated load, estimated by the simply-supported beam formula $\delta = PL^3/(48EI)$, would exceed 50 metres — a number that demonstrates a 1,000-metre girder without cable support is utterly inconceivable.

With cable tension participating, the midspan deflection calculated from the tension-bending beam model is approximately 0.08 m. Geometric stiffness compresses deflection by several hundred times. This is the fundamental reason the stiffening girder section can be so small: it is not the primary provider of stiffness.

This comparison reveals the stiffness composition of suspension bridges: for long-span suspension bridges, the contribution of cable geometric stiffness to total vertical stiffness far exceeds that of the stiffening girder's elastic stiffness.

2.3 Three Remarkable Conclusions

The greatest value of the equivalent beam model is not that it gives us another calculation method, but that it reveals three conclusions that defy intuition yet cannot be refuted.

Conclusion 1: Deflection is independent of section size

When geometric stiffness greatly exceeds elastic stiffness — the usual case for long-span suspension bridges — deflection is governed primarily by cable tension H. Through simplified derivation, a remarkably concise expression for deflection-to-span ratio emerges:

$$\delta/L \approx (5/48) \cdot (p/g) \cdot \alpha$$

where p is live load, g is dead load, and α = f/L is the sag-to-span ratio.

Examine the right-hand side carefully. p/g is the live-to-dead load ratio; α is the sag-to-span ratio. The formula contains no EI, no sectional area A, no section parameter of any kind.

This is the most striking conclusion of this chapter: the deflection of a suspension bridge is independent of the stiffening girder section size. Enlarging the girder section has almost no effect on deflection control. What matters is the sag-to-span ratio and the live-to-dead load ratio.

The comparison with Chapter 1 is particularly telling. The arch bridge deflection formula contains no moment of inertia I, but at least includes sectional area A — controlling arch deflection depends on EA. The suspension bridge is more extreme still: not even a section parameter appears. This is geometric stiffness overwhelming section stiffness in its purest form.

Numerical check: for $\alpha = 1/10$ and $p/g = 1/5$, $\delta/L \approx (5/48) \times (1/5) \times (1/10) \approx 1/480$. For a 1,000-metre span, maximum deflection ≈ 2 m — consistent with practice.

Conclusion 2: Natural frequency is also independent of section size

When geometric stiffness dominates, the fundamental vertical bending frequency of a suspension bridge approximates a string vibration formula:

$$f_1 \approx 0.55 / \sqrt{(\alpha \cdot L)}$$

This result is equally striking: the fundamental vertical bending frequency depends only on span and sag-to-span ratio, and is independent of the stiffening girder section and material.

The reason: geometric stiffness is produced by dead load, and vibration inertia is also produced by dead load (mass); both arise from the same source and cancel in the frequency expression. This is identical in principle to the pendulum period depending only on pendulum length and being independent of bob mass — a shared characteristic of all gravity-dominated structures.

Note that the above formula is based on a simplified model in which geometric stiffness completely dominates. In real bridges, the stiffening girder's flexural stiffness and boundary constraints raise the actual fundamental frequency to 1.2–1.5 times the theoretical lower bound. But the core conclusion stands: the controlling factors for frequency are α and L, not section properties. For long-span suspension bridges, one cannot expect to improve frequency or dynamic performance by adding weight to the stiffening girder. The only effective means are two: shorten the span, or reduce the sag-to-span ratio.

Conclusion 3: Dead load is the source of stiffness

This is the most counter-intuitive conclusion of all.

In girder bridges, dead load is a pure burden — greater weight means larger bending moment, greater deflection, and a necessarily larger section. In arch bridges, dead load is also the enemy — greater weight means greater axial compression and higher instability risk. But in suspension bridges, dead load plays an entirely different role: greater dead load produces greater cable tension, higher geometric stiffness, and a stiffer bridge.

This can be read directly from the deflection-to-span formula: the ratio is proportional to p/g. Dead load g appears in the denominator; the larger g, the smaller the deflection-to-span ratio. Increasing dead load is increasing stiffness.

This explains a phenomenon familiar to engineers but perhaps not deeply considered: why must the stiffening girder of a railway suspension bridge be far heavier than that of a highway bridge? Not because strength is insufficient — but because railway live loads are large, and a greater dead load is needed to suppress the live-to-dead load ratio and thereby control deflection. The essential effect of increasing girder weight is to raise g (and thus stiffness), not to strengthen EI.

In the language of Chapter 1: in an arch bridge, self-weight is the accomplice of negative geometric stiffness; in a suspension bridge, self-weight is the source of positive geometric stiffness. The same gravity, under different signs of geometric stiffness, plays completely opposite roles — yet another manifestation of the arch–cable mirror symmetry introduced at the end of Chapter 1.

2.4 The True Role of the Stiffening Girder: Enabler, Not Primary Carrier

If the stiffening girder's flexural stiffness contributes so little to overall deflection and frequency, what exactly is its role? Why not simply remove it?

The answer: the stiffening girder exists not to provide stiffness, but to ensure that the cable's geometric stiffness can effectively serve the bridge deck. Its role can be summarised in three aspects.

First, load distribution. The cable is a flexible member that cannot independently resist local deformation caused by concentrated loads. The stiffening girder distributes concentrated loads longitudinally so that cable forces are more uniform. The stiffer the girder, the wider the distribution; the more flexible the girder, the more concentrated the cable forces and the larger the local deformations.

Second, local deformation control. Even when the overall deflection-to-span ratio is satisfied, asymmetric loading (e.g. half-span live load) can produce large local rotations and curvatures near midspan. These local indicators are governed primarily by the girder's EI — an important constraint on girder section design.

Third, torsional stiffness. This is often the true governing criterion for stiffening girder section design in long-span suspension bridges. Flutter stability depends on the relationship between vertical bending and torsional frequencies — the wider the separation, the higher the critical wind speed.

Vertical bending frequency is controlled primarily by cable tension (geometric stiffness); torsional frequency is controlled primarily by the girder's torsional stiffness GJ and the distance between cable planes. Flutter stability is, in essence, a matching problem between two types of stiffness.

The Tacoma Narrows Bridge, opened in 1940, collapsed in wind after only four months. The lesson was not insufficient strength but precisely a stiffness-matching failure: the solid-web plate girder had large EI but grossly inadequate GJ; vertical bending and torsional frequencies were too close together, leading to wind-induced coupled vibration. This was not a strength problem but a pure stiffness-matching problem.

Modern long-span suspension bridges universally adopt streamlined closed steel box girders. A core consideration is that their torsional stiffness is far superior to that of truss girders and I-girders, enabling the separation between vertical bending and torsional frequencies to be widened.

In one sentence: the stiffening girder is not the "load-carrying member" of a suspension bridge but its "enabling member" — it ensures that cable geometric stiffness can be effectively transmitted to the bridge deck, and provides the torsional stiffness necessary for flutter stability. This relationship — "geometric stiffness requires elastic stiffness to maintain it" — is the core characteristic of the suspension bridge stiffness system.

The comparison with Chapter 1 is worth noting. In an arch bridge, EI is the "suppressor" of negative geometric stiffness — maintaining the compression state without buckling. In a suspension bridge, EI and GJ are the "transmitters" of positive geometric stiffness — delivering the cable's restoring force to the deck. One is "defence", the other "transmission", but the essence is the same: section stiffness is the enabling condition for geometric stiffness. Without this enabling condition, no amount of geometric stiffness can be converted into usable deck stiffness.

2.5 Sag-to-Span Ratio: The Core Parameter

Consolidating the preceding analysis, the parameter causal chain of a suspension bridge can be stated concisely:

Sag-to-span ratio α → Cable tension H → Geometric stiffness → {Deflection, Frequency, Stability} + Minimum requirements for stiffening girder

Once the sag-to-span ratio is fixed, the order of magnitude of essentially all core indicators is determined. The choice of stiffening girder section is a subordinate decision made on this basis — subordinate to the torsional demand of flutter stability, not to the flexural demand of vertical stiffness.

Trend comparison for a 1,000-metre highway suspension bridge (dead load 200 kN/m, live load 40 kN/m):

Criterion	α = 1/8	α = 1/10	α = 1/12
Cable sag (m)	125	100	83
Cable tension H (kN)	250,000	312,500	375,000
Deflection-to-span (est.)	≈ 1/400	≈ 1/500	≈ 1/600
Cable quantity ratio	0.80	1.00	1.20

From α = 1/8 to 1/12, the deflection-to-span ratio improves from approximately 1/400 to 1/600, but cable quantity increases by roughly 50 per cent. This is the core trade-off of suspension bridge design: purchasing higher stiffness with more material.

The truly compelling evidence is the collective choice of engineering practice:

Bridge	Main span (m)	Actual α	Completion
Akashi Kaikyō Bridge, Japan	1,991	≈ 1/10.0	1998
Xihoumen Bridge, China	1,650	≈ 1/10.8	2009
Runyang Bridge, China	1,490	≈ 1/10.2	2005
Tsing Ma Bridge, Hong Kong	1,377	≈ 1/9.4	1997
Golden Gate Bridge, USA	1,280	≈ 1/9.3	1937

Five bridges, spanning nearly ninety years, from different countries and design teams: sag-to-span ratios cluster without exception between 1/9 and 1/11. This is neither coincidence nor code requirement, but the collective confirmation of the stiffness-versus-material trade-off predicted by theory.

One implicit premise is worth making explicit: the preceding analysis assumes immovable anchorages. The condition for H to be constant is that the anchorages undergo no horizontal displacement. If anchorage horizontal stiffness is limited, the effective cable tension decreases and geometric stiffness falls. This is why the anchorages of long-span suspension bridges must be founded on bedrock — the influence of boundary conditions on stiffness will be discussed systematically in Chapter 6.

2.6 The Suspension Bridge’s Place on the Stiffness Spectrum

Placing the suspension bridge alongside the continuous girder bridge makes the contrast clearer:

Feature	Continuous Girder Bridge	Suspension Bridge
Primary stiffness source	Section flexural stiffness EI	Geometric stiffness from cable tension
Means of raising stiffness	Enlarge section, shorten span, add supports	Reduce sag-to-span ratio, increase dead load
Role of dead load	Pure burden (heavier = worse)	Source of stiffness (heavier = stiffer)
Typical depth/span	1/20 to 1/50	1/200 to 1/800 (stiffening girder)
Core design parameter	Section, span, support conditions	Sag-to-span ratio (the core)
Critical stability mode	Construction-stage concerns	Wind-induced torsion (vertical self-stabilising)

The left column is the familiar "elastic stiffness world" — stiffness depends on section, dead load is burden. The right column is the suspension bridge's "geometric stiffness world" — stiffness depends on geometric configuration, dead load is the source of stiffness. Two worlds with entirely different logics, but both governed by stiffness.

Returning to the stiffness spectrum established in the Introduction: the suspension bridge occupies a position precisely symmetric to the arch bridge — at the extreme end of the positive geometric stiffness region. The arch bridge, in the $K_G < 0$ region, offers "risky stiffness". The suspension bridge, in the $K_G > 0$ region, offers "living stiffness": the greater the load, the more stable. But this "living" quality also implies vulnerability — if cable tension is lost (e.g. cable rupture), geometric stiffness vanishes instantly, unlike section stiffness which, as an intrinsic material property, is always present.

The suspension bridge also reveals a unique physical phenomenon: a "positive feedback" of vertical stability. In an arch bridge, axial compression produces negative geometric stiffness — greater load drives the structure closer to instability, a dangerous negative feedback. In a suspension bridge, greater vertical load produces greater cable tension, higher geometric stiffness, and greater stability — a beneficial positive feedback. This does not mean suspension bridges have no stability problems — their danger lies not in the vertical direction but in wind-induced torsion. But the vertical self-stabilising mechanism is indeed one of the mechanical foundations enabling suspension bridges to achieve ultra-long spans.

The two share a deeper commonality: in neither case does stiffness truly originate from the section. The arch, through its form, converts bending into axial compression, relying on EA in place of EI. The suspension bridge goes further — depending almost entirely on geometric stiffness, with EI serving merely as the enabling condition. From girder to arch to suspension bridge, the centre of gravity of stiffness shifts progressively from “section” to “form”. This is the main thread of stiffness spectrum evolution.

2.7 What the Equivalent Beam Teaches Us

The stiffness logic of suspension bridges can be summarised in three statements.

First, the vertical stiffness and natural frequency of a suspension bridge are governed primarily by cable tension (geometric stiffness); the contribution of the stiffening girder’s flexural stiffness is small. This explains why the girder can be far shallower than a continuous girder of the same span, and why enlarging the girder section brings almost no improvement in deflection.

Second, the sag-to-span ratio is the sole core parameter of a suspension bridge. It determines geometric stiffness level through cable tension, and thereby determines deflection, frequency, and nearly all core indicators. Its concentration between 1/9 and 1/11 reflects the trade-off between stiffness demand and cable economy.

Third, the stiffening girder’s role is “enabling” rather than “primary carrying”. The relationship — “geometric stiffness requires elastic stiffness to maintain it” — is the core characteristic of the suspension bridge stiffness system.

The most important formula in this chapter is the deflection-to-span ratio: $\delta/L \approx (5/48)(p/g)\alpha$. Its striking feature is that deflection-to-span ratio is entirely independent of the stiffening girder section. If Chapter 1’s arch formula proved that “controlling arch deflection depends on EA, not EI”, then the suspension bridge formula goes further, proving that “controlling suspension bridge deflection does not even depend on the section itself”.

A final word on the value of simplified models. The equivalent beam formulae of this chapter cannot replace finite-element analysis. But they do what finite elements cannot: establish causal relationships among parameters. Finite elements will tell you a suspension bridge has a fundamental frequency of 0.35 Hz, but they will not tell you what that frequency principally depends on, which parameter is most effective to adjust, or why certain adjustments are virtually ineffective. The equivalent beam model provides not computational power but judgement. At the concept-design stage, knowing which parameter to adjust — and in which direction — may be more important than calculating precisely how much to adjust it.

The next chapter turns to cable-stayed bridges. If the suspension bridge is a structure dominated by geometric stiffness — with the sag-to-span ratio as virtually the only parameter, leaving the designer a narrow range of choice — then the cable-stayed bridge will present an entirely different picture: multiple independently adjustable stiffness parameters, stiffness no longer passively received but actively planned. This is the critical step from “passive stiffness” to “active stiffness”.

Chapter 3

Cable-Stayed Bridges: The Modern Structure
That Actively Configures Stiffness

Why cable-stayed bridges are "truly modern bridges"

3.1 From "Integrated" to "Assembled"

The preceding chapter used Li Guohao's equivalent beam model to reduce the stiffness essence of suspension bridges to a concise picture: geometric stiffness from cable tension dominates vertical stiffness, natural frequency, and vertical stability, while the sag-to-span ratio is the single core parameter controlling all of it. The conclusion: the design freedom of a suspension bridge is quite limited — once the sag-to-span ratio is fixed, the stiffness landscape is essentially set.

The situation with cable-stayed bridges is fundamentally different. Although also a cable-supported bridge, the cable-stayed bridge possesses three characteristics that the suspension bridge lacks. First, the force in each stay can be controlled independently, rather than being dictated uniformly by cable geometry. Second, the structure is highly statically indeterminate, with internal force distribution depending on member stiffness ratios rather than equilibrium alone. Third, elastic stiffness and geometric stiffness are strongly coupled, with multiple parameters independently adjustable.

These characteristics mean that cable-stayed bridge stiffness is not passively received but actively planned. In the language of the Introduction: suspension bridge design searches for the optimum within a virtually one-dimensional parameter space (sag-to-span ratio), while cable-stayed bridge design actively configures the stiffness distribution across a multi-dimensional parameter space. This is perhaps the mechanical justification for calling cable-stayed bridges "truly modern bridges" — their modernity lies not in the era of their appearance, but in their embodiment of the core idea of modern structural design: actively configuring stiffness rather than passively checking it.

3.2 Vertical Support Stiffness of a Single Stay: The Origin of $\sin^3\theta$

Each stay cable of a cable-stayed bridge acts as an elastic bearing — when the deck moves down slightly, the stay produces an upward restoring force. The magnitude of this restoring force depends on the vertical support stiffness k of that stay.

To understand why this stiffness is proportional to $\sin^3\theta$, the key is to ask: through how many steps is a unit vertical deck displacement finally converted into a vertical restoring force? The answer is three steps, each “losing” one factor of $\sin\theta$.

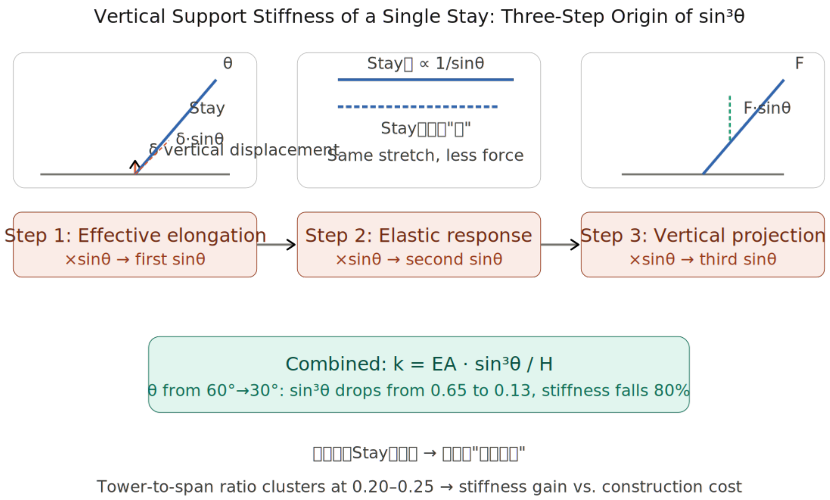

Figure 4 — Three-step origin of $\sin^3\theta$: how a unit vertical deck displacement is converted into a vertical restoring force through three successive reductions, each costing one factor of $\sin\theta$.

Step 1 — Effective elongation. The deck moves vertically downward, but the stay is inclined. Only the component along the cable axis actually elongates it: the effective component is the vertical displacement multiplied by $\sin\theta$. The smaller θ (the flatter the stay), the smaller the proportion that actually stretches the cable. First $\sin\theta$.

Step 2 — Elastic response. For the same elongation, a longer cable produces less elastic force — just as compressing a soft spring by 1 mm gives less force than compressing a stiff spring. A flatter cable at the same tower height is longer and therefore “softer”. Cable length is proportional to $1/\sin\theta$. Second $\sin\theta$.

Step 3 — Vertical projection. The cable force acts along the cable axis. To obtain the vertical restoring force actually applied to the deck, one must take the vertical component, multiplying by $\sin\theta$ once more. Third $\sin\theta$.

Three steps combined, vertical support stiffness is proportional to $\sin^3\theta$. Including cable material and section (elastic modulus E, sectional area A) and the relationship between cable length and tower height H:

$$k = EA \cdot \sin^3\theta / H$$

Three important features follow naturally from the three steps.

First, the cubic relationship means cable angle is extremely sensitive to vertical stiffness. When θ drops from 60° to 30°, $\sin^3\theta$ falls from 0.65 to 0.13 —

vertical stiffness drops by approximately 80 per cent. This explains why far-field mid-span stays are far less effective than near-tower stays: mid-span cable angles are smallest, creating the natural “stiffness trough” of cable-stayed bridges.

Second, the taller the tower, the greater the average vertical stiffness of the stays, but with diminishing returns. This explains why tower-to-span ratios cluster at 0.20–0.25: beyond this, stiffness gains flatten while construction costs continue to accumulate.

Third, stay stiffness is proportional to cable axial stiffness EA. This is an important distinction from suspension bridges — the designer can directly adjust support stiffness by choosing the cable section. In a suspension bridge, cable tension is dictated by dead load and sag-to-span ratio; the designer cannot freely choose it.

For small and medium-span cable-stayed bridges (main span below 500 m), stays are relatively short and the formula applies directly. When the main span exceeds 800 m, the longest stays can reach over 500 m, and the sag effect of cable self-weight significantly reduces the effective elastic modulus (Ernst effect) — one of the fundamental reasons for the ceiling on cable-stayed bridge spanning capability.

3.3 The Elastic Foundation Beam Model: The Essential Distinction from Suspension Bridges

Modern cable-stayed bridges typically have stay spacings of 8–15 m and large numbers of stays. When the spacing is much less than the structural characteristic length, the discrete stay supports can be replaced by a continuous elastic foundation. The governing equation of the main girder then becomes the classical Winkler elastic foundation beam equation.

To appreciate the essential difference, imagine two physical scenes before examining the equations.

Scene one: fix a rope at both ends and pull it taut; press down at a point in the middle. The resistance comes from the entire rope “sensing” the bent shape and generating a tension component — press one point, and the full beam participates in the response. This is exactly the logic of a suspension bridge cable: geometric stiffness couples with the curvature of the stiffening girder (second derivative of displacement) — a global constraint.

Scene two: place a beam on a row of independent springs and apply a load at one point. Each spring senses only the displacement at its own location, with no connection to adjacent springs. This is exactly the logic of cable-stayed bridge stays — each “sees” only the displacement at its own anchorage point, responding independently, without interfering with other stays.

	Suspension Bridge (tension-bending beam)	Cable-Stayed Bridge (elastic foundation beam)
Constraint type	Cable tension × curvature (2nd derivative)	Stay × displacement (0th derivative)
Constraint extent	Integrated: tension uniformly affects entire beam	Assembled: each stay constrains local region
Stiffness adjustability	Limited: cable geometry hard to change after setting	Rich: each stay independently adjustable

In short: suspension bridge stiffness is “integrated” — change one thing, change everything. Cable-stayed bridge stiffness is “assembled” — adjustable point by point. This is the mathematical root of “active stiffness”.

Characteristic length: the scale of adequate stay support

Elastic foundation beam theory contains an important characteristic parameter $\lambda = (4EI/k)^{1/4}$. Its physical meaning: the range over which deflection propagates significantly under a local load. When λ is much smaller than the main span L, stay support effectively constrains deformation; when λ approaches L, stay support is insufficient and the main girder degenerates into an ordinary continuous beam. The basic design requirement for cable-stayed bridges is therefore $\lambda/L \ll 1$.

In practice, λ/L for cable-stayed bridges typically falls in the range 0.02–0.10 — well below 1, confirming the dominance of stay support. This parameter is also a direct basis for judging whether stay spacing is reasonable and whether the match between girder stiffness and stay stiffness is appropriate.

3.4 Core Control Parameters

Vertical deflection

When stay support stiffness dominates ($\lambda/L \ll 1$), vertical deflection approximates the ratio of live load to stay support stiffness. Substituting the $\sin^3\theta$ formula and accounting for stay spacing s, the deflection-to-span ratio can be expressed as:

$$\delta/L \approx (p \cdot s \cdot H) / (EA \cdot \sin^3\theta \cdot L)$$

Effective means to reduce this ratio: decrease stay spacing s (increase stay count); increase cable EA; increase tower height (raise $\sin\theta$); reduce live load p. Increasing the main girder section has limited effect — similar to the suspension bridge conclusion that “enlarging the stiffening girder section has almost no effect on deflection”, but with a different mechanism: in suspension bridges geometric stiffness dominates; in cable-stayed bridges stay support stiffness dominates.

Notably, cable-stayed bridges possess an adjustment lever that suspension bridges lack: stay spacing s. Reducing spacing means increasing “support density”, directly raising system stiffness. This is a design degree of freedom unique to cable-stayed bridges.

Fundamental frequency

The fundamental vertical bending frequency of the elastic foundation beam is approximately $f_1 \approx (1/2\pi)\cdot\sqrt{(\bar{k}/m)}$. When stay support stiffness greatly exceeds main girder flexural stiffness, the frequency is governed primarily by the ratio of stay support stiffness to mass, with little dependence on girder EI.

But cable-stayed bridges have an important frequency characteristic absent in suspension bridges: low-order modes are dominated by stay support stiffness, while high-order modes are dominated by main girder stiffness. The reason: in the elastic foundation beam, the girder stiffness term in higher-order frequencies contains an n^4 factor that grows rapidly with mode number, soon exceeding the stay support stiffness term. This means main girder stiffness plays a significant role in high-order modes and cannot be neglected.

Flutter stability

In flutter stability, cable-stayed bridges hold an important advantage over suspension bridges. Vertical bending frequency is controlled primarily by stay support stiffness — adjustable through cable section, spacing, and tower height. Torsional frequency is controlled primarily by the main girder’s torsional stiffness GJ and the distance between cable planes. The designer can adjust vertical and torsional stiffness relatively independently, thereby actively optimising the ratio of vertical bending to torsional frequency. This “stiffness-decoupled control” capability is something suspension bridges lack.

3.5 The Three Roles of the Tower

The tower of a cable-stayed bridge is more than an anchorage point for stays; it affects stiffness at three levels.

First, geometric enhancement. Tower height determines stay inclination; inclination determines vertical stiffness through $\sin^3\theta$. This is the most direct pathway from tower height to stiffness.

Second, force-flow conversion. The tower converts the horizontal component of stay forces into vertical foundation pressure. In self-anchored systems, it also transmits axial compression to the main girder, producing P-Δ effects that partially offset the positive stiffness provided by stays.

Third, tower stiffness safeguard. The tower’s longitudinal stiffness directly affects tower-top horizontal displacement. Under asymmetric loading between main span and side spans, tower-top sway simultaneously alters the geometric relationships of all stays, causing support stiffness to decline in unison. This is a global stiffness degradation mechanism — insufficient tower stiffness triggers a system-wide stiffness drop. In ultra-long-span cable-stayed bridges (main span above 800 m), tower sections are often governed not by strength but by longitudinal stiffness.

Thus, the tower-to-span ratio in cable-stayed bridges plays a role analogous to the sag-to-span ratio in suspension bridges — both are core form parameters controlling stiffness efficiency. Tower-to-span ratios cluster at 0.20–0.25, just as suspension bridge sag-to-span ratios cluster at 1/9–1/11 — both the result of balancing stiffness supply against structural cost.

3.6 Stay Force Adjustment: Active Stiffness Configuration

Cable-stayed bridges possess a capability entirely absent in suspension bridges: actively controlling stiffness distribution by adjusting stay forces.

A cable-stayed bridge is a highly statically indeterminate structure; a typical twin-tower cable-stayed bridge may have 80–200 degrees of indeterminacy. In statically indeterminate structures, internal force distribution depends on member stiffness ratios — stiffer members attract more load, more flexible members shed it. Adjusting stay forces is therefore, in essence, adjusting stiffness distribution — the core mechanical content of cable-stayed bridge construction control.

The so-called “ideal completed state” can be redefined in the language of stiffness: it is the state in which the distribution of stay support stiffness is optimally matched to the dead-load distribution, minimising main-girder bending moments under dead load. In a suspension bridge this matching is achieved automatically by cable geometry (a parabola naturally balances uniform dead load); in a cable-stayed bridge it must be achieved actively through tensioning adjustments — a direct manifestation of the “designability” of cable-stayed bridge stiffness.

High static indeterminacy also brings another advantage: stiffness redundancy. Even if a stay loses stiffness through fatigue or corrosion, adjacent stays and the main girder can redistribute loads; the structure will not collapse from a single stay failure. Moreover, stays are replaceable members — this “stiffness maintainability” is difficult for suspension bridges or arch bridges to match.

3.7 Multi-Parameter Adjustability: The Cable-Stayed Design Space

Consolidating the analysis above, the most prominent distinction between cable-stayed and suspension bridges is the number and independence of adjustable parameters.

A suspension bridge’s deflection, frequency, and stability are essentially all determined by one parameter: the sag-to-span ratio. A cable-stayed bridge has stay spacing, tower height, cable section, main girder stiffness, and stay force distribution — multiple independently adjustable parameters. Taking a 500-metre main-span twin-tower cable-stayed bridge as example:

Stay spacing: reducing from 12 m to 8 m raises support stiffness by approximately 50 per cent and fundamental frequency by approximately 22 per cent — but increases total cable quantity by approximately 48 per cent. This is the trade-off between support density and material economy.

Tower-to-span ratio: increasing from 0.18 to 0.26 raises support stiffness by approximately 75 per cent and fundamental frequency by approximately 32 per cent — but tower height rises from 90 m to 130 m, with significantly greater construction difficulty and stability concerns.

Each parameter has its own stiffness contribution efficiency and economic cost; the designer can search for the optimal combination within this multi-dimensional space. This "multi-parameter adjustability" is why cable-stayed bridges can flexibly adapt to a wide range of spans and site conditions, and the fundamental reason they are highly competitive across the broad span range of 200–1,200 metres.

3.8 The Cable-Stayed Bridge's Place on the Stiffness Spectrum

Returning to the stiffness spectrum of the Introduction: the cable-stayed bridge occupies a unique position — it spans both positive and negative geometric stiffness regions and possesses the greatest degree of design freedom.

Compared with the arch and suspension bridges, the distinctive feature of the cable-stayed bridge is that it does not operate at a single point on the stiffness spectrum but across a region. By adjusting multiple parameters, the designer can actively choose the optimal stiffness configuration within this region. Arch and suspension bridge design resembles searching for the best position along a narrow path; cable-stayed bridge design is more like planning a layout across open terrain.

Stage	Bridge type	Mode of stiffness configuration
Passive	Girder bridge	Section determines everything; no geometric help
Static	Arch bridge	Form fixed once set; EA dominates
Single-parameter	Suspension bridge	Sag-to-span ratio is essentially the only parameter
Multi-parameter active	Cable-stayed bridge	Spacing, tower height, section, force — all adjustable

This is the evolutionary trajectory from "passive stiffness" to "active stiffness".

A systematic comparison of core characteristics makes the differences even more vivid:

Feature	Suspension Bridge	Cable-Stayed Bridge
Dominant stiffness mechanism	Cable geometric stiffness (integrated)	Stay support elastic stiffness (assembled)
Governing equation	Tension-bending beam: tension × curvature	Elastic foundation beam: spring × displacement
Core form parameters	Sag-to-span ratio (virtually the only one)	Spacing, tower height, cable section (multiple)
Role of dead load	Source of stiffness (heavier = stiffer)	Burden (compensated by stay forces)
Stay force adjustability	Not adjustable after cable geometry is set	Each stay force independently adjustable
Design philosophy	"How much dead weight to exchange for stiffness?"	"How to actively configure stiffness?"

The final row captures the essential distinction. Suspension bridge design is fundamentally about answering "how much self-weight to trade for how much stiffness" — a narrow range of choice. Cable-stayed bridge design is about answering "how to actively configure stiffness" — a broad range of choice.

3.9 From Passive Stiffness to Active Stiffness

The stiffness logic of cable-stayed bridges can be summarised in three statements.

First, the vertical support stiffness of a stay is proportional to the cube of the cable inclination angle — the $\sin^3\theta$ relationship makes stiffness extremely sensitive to cable angle. This explains why far-field mid-span stays are far less effective than near-tower stays, and is the root cause of the "stiffness trough" in cable-stayed bridges.

Second, the stiffness system of a cable-stayed bridge is "assembled" — unlike the "integrated" system of a suspension bridge, each stay is an independently controllable elastic support whose stiffness, position, number, and pretension can all be actively set by the designer.

Third, cable-stayed bridges possess multiple independently adjustable stiffness parameters. This multi-parameter adjustability makes active stiffness configuration possible — the core characteristic distinguishing cable-stayed bridges from all other bridge types.

From the suspension bridge's "tension-bending beam" to the cable-stayed bridge's "elastic foundation beam", these two chapters have used the same approach — reducing cable-supported structures to their most basic beam models — to unify our stiffness understanding of the two cable bridge types. The difference between the two governing equations — minus sign versus plus sign, second derivative versus zeroth order, integrated versus assembled — precisely corresponds to the essential distinction in stiffness configuration between the two bridge types.

The cable-stayed bridge marks the transition from passive to active stiffness configuration. This transition is not merely technical progress but a significant conceptual leap: the structure is no longer the passive bearer of stiffness but its active controller.

As with the preceding chapters, the simplified models here cannot replace finite-element analysis. They do what finite elements cannot: reveal causal relationships among parameters. When finite elements tell you a cable-stayed bridge has a fundamental frequency of 0.35 Hz, they will not tell you whether that frequency is governed primarily by stay support stiffness or by girder stiffness, whether adjusting spacing or girder section is more effective, or which other indicators a change in tower height will affect. The elastic foundation beam model provides precisely this kind of judgement.

The next chapter returns to the starting point of the stiffness spectrum — the girder bridge. If arch, suspension, and cable-stayed bridges all represent different modes of breaking through the girder bridge's L^4 law, then understanding the L^4 law itself — how it constrains the girder bridge, and how the girder bridge strives to extract every last bit of potential within this constraint — is the foundation for understanding the entire stiffness spectrum.

Chapter 4

Girder Bridges: The Reference Frame for All Systems

The origin of the stiffness spectrum and the boundary of spanning capability

4.1 The Overlooked Simplicity

In the preceding three chapters we examined the geometric reconstitution of arch bridges, the geometric-stiffness dominance of suspension bridges, and the active stiffness configuration of cable-stayed bridges. But one question has been embedded throughout without separate scrutiny: what is the stiffness logic of the girder bridge?

The question seems too simple. Girder stiffness is EI — two letters and you are done. But it is precisely this apparent simplicity that conceals questions worth pondering. Why does a girder bridge have a hard upper limit on span? Why can a continuous girder span further than a simply-supported one? Does prestress increase strength or stiffness? What is the mechanical logic of a variable-depth girder? Most importantly: what essential distinction in stiffness logic between the girder and the arch, suspension, and cable-stayed bridge makes the invention of those other systems necessary?

Answering these questions is not merely a matter of completing Part One. The girder bridge is the reference frame for all structural systems. In the preceding chapters, we modelled the suspension bridge as a “tension-bending beam” and the cable-stayed bridge as an “elastic-foundation beam”. Why does every model return to the beam? Because the beam is the purest form of stiffness — only elastic stiffness K_E, with no participation of geometric stiffness K_G. Every other system is, in essence, the beam baseline with an additional stiffness mechanism superimposed. Without understanding the beam, one cannot truly understand what is innovative about the others.

4.2 The L^4 Law: The Brutal Fourth Power

The midspan deflection of a simply-supported beam under uniformly distributed load p:

$$\delta = 5pL^4 / (384EI)$$

Simple as it is, this formula contains the entire essence of girder bridge stiffness logic. The L^4 in the numerator is decisive: deflection is proportional to the fourth power of span and inversely proportional to section stiffness. This is the “L^4 law” of girder bridges.

The implication of L^4 is brutal. Double the span: deflection increases sixteenfold. To maintain the same deflection-to-span ratio, section stiffness must increase sixteenfold. For a prismatic beam, moment of inertia I is roughly proportional to the cube of depth h, so depth must increase to approximately 2.5 times — but self-weight is also rising rapidly. A vicious cycle emerges: greater span → higher stiffness demand → larger section → heavier self-weight → greater deflection → need for still larger section. At some span this cycle ceases to converge — that is the span limit of girder bridges.

It is precisely the existence of the L^4 law that makes the arch bridge's L^2/EA (Chapter 1) and the suspension bridge's "deflection independent of section" (Chapter 2) such remarkable breakthroughs. They are not marginal improvements within the L^4 framework but reductions in span exponent achieved by changing the mechanical load path.

Depth-to-span ratio: the sole control parameter

Rearranging slightly: for a box section, the moment of inertia can be written $I \approx \alpha b h^3$, where α is a section efficiency coefficient. Substituting:

$$\delta/L \;\propto\; (p/Eb) \cdot (L/h)^3$$

Deflection-to-span ratio is proportional to the cube of the span-to-depth ratio L/h. Once material and section type are fixed, stiffness is almost entirely determined by L/h. This explains why span-to-depth ratios cluster in narrow bands: highway simply-supported girders typically at L/h = 15–20, continuous girders at 18–25. The cubic relationship means increasing L/h by 30 per cent increases deflection by approximately 120 per cent — the adjustment margin is extremely limited.

Frequency: the same logic in another guise

The fundamental vertical bending frequency of a simply-supported beam is $f_1 = (\pi^2/2L^2)\cdot\sqrt{(EI/m)}$. The L^2 likewise originates from the L^4 law (frequency is the square root of stiffness). Deflection and frequency are governed by the same parameter — span-to-depth ratio — with no independent degree of freedom for adjustment. This is a marked contrast with cable-stayed bridges, where stay forces can independently influence vertical bending and torsional stiffness.

4.3 Four Strategies Within the L^4 Law

Girder bridge engineers are not helpless before the L^4 law. Over its long development history, four strategies have been devised to extend spanning capability. Each is ingenious, but all share one feature: none breaks the L^4 law — they maximise effort within it.

Continuous girder: trading boundary conditions for stiffness efficiency

A simply-supported beam has a midspan moment of $pL^2/8$. Making adjacent spans continuous introduces negative moments at supports and reduces the midspan positive moment to roughly 40–60 per cent of the simply-supported value.

From a stiffness perspective, the essence of the continuous girder is: through a change in boundary conditions, bending moment is redistributed within the span. The L^4 law itself is unchanged, but the “effective span” is reduced — the redistribution of moment is equivalent to converting one large-span problem into several smaller-span problems. The penalty is large negative moments at supports, imposing new demands on sections and bearings.

Continuity can extend the practical span of girder bridges by roughly 30–50 per cent. Significant, but not revolutionary — the mode of stiffness generation is unchanged; only its utilisation efficiency is improved.

One premise deserves emphasis: the benefit of continuity requires rigid supports. If intermediate supports have limited vertical stiffness (e.g. tall piers with flexible columns and foundations), differential settlement induces parasitic moments that may erode the advantage entirely. This is why mountain bridges with tall piers often adopt simply-supported spans — not because designers are unaware of continuity’s benefits, but because boundary conditions (pier flexibility) prevent those benefits from being realised. This will be discussed in detail in Chapter 6.

Variable depth: matching the EI distribution to the moment

A prismatic beam is obviously irrational: the section is uniform everywhere, yet the bending moment distribution is highly non-uniform. Varying the depth along the span — deeper where moments are large, shallower where they are small — significantly improves material utilisation.

From a stiffness perspective, the essence of the variable-depth girder is: matching the distribution of EI to the distribution of bending moment. The typical continuous-beam practice — deepened at supports (large moments), thinned at midspan (small moments) — gives the familiar “haunch” profile. This shape is not an aesthetic choice but a mechanical necessity of stiffness matching.

Prestress: not adding stiffness but preventing its loss

Prestress is usually understood as a strength measure — using precompression to counteract tensile stress and prevent cracking. This is correct but incomplete. From a stiffness perspective, the essential role of prestress is to maintain the effective EI of the section.

This requires explanation. When a concrete beam cracks in tension, the moment of inertia of the cracked section drops sharply — typically to 30–60 per cent of the uncracked value. Cracking is therefore not only a strength issue but a stiffness issue. Once cracked, the effective EI drops dramatically and structural stiffness falls with it.

Prestress, by preventing cracking (or controlling crack width), maintains the section's effective stiffness. For long-span girder bridges, this role is decisive. Without prestress, the actual stiffness of a concrete beam might be only half its nominal value or less; large spans would be out of the question.

But an important distinction must be drawn: prestress does not add stiffness — it merely prevents stiffness from being lost. It is a "stiffness maintenance" measure, not a "stiffness creation" measure. The girder's stiffness still comes from the section's EI; prestress merely ensures that EI can be fully realised. This is entirely different from the cable tension of a suspension bridge, which is a direct source of stiffness.

Composite section: superimposing the EI-to-weight ratios of two materials

A steel–concrete composite girder places the concrete deck slab in the compression zone, connected to the steel girder through shear connectors. The composite moment of inertia far exceeds that of the steel girder alone — typically 1.5–2.5 times.

The distinction from prestress is fundamental. Prestress "maintains EI" — preventing cracking-induced stiffness loss. The composite section "superimposes EI" — combining the strengths of two materials to achieve section efficiency unattainable by either alone.

The composite section also offers an important stiffness advantage: the far higher strength-to-weight ratio of steel makes the composite girder much lighter than a pure concrete girder of equal stiffness. Reduced self-weight directly lowers the self-weight "consumption" of stiffness, effectively raising the span ceiling. This explains why the world's longest-span girder bridges almost all employ steel or composite sections.

Common features of the four strategies

Strategy	Means	Typical effect	Essence
Continuous girder	Redistributes bending moment	Span increase 30–50%	Boundary condition optimisation
Variable depth	EI distribution matches moment	Significant material utilisation improvement	Spatial optimisation of EI
Prestress	Prevents cracking; maintains effective EI	Stiffness realisation rate doubles or more	Maintenance of stiffness
Composite section	Superimposes contributions of two materials	EI/weight ratio significantly improved	Superposition of stiffness

For the concrete path, the first three strategies combined yield the standard configuration of modern long-span girder bridges: prestressed concrete continuous variable-depth box girder. For the steel and composite path, the composite section leverages a higher EI-to-weight ratio to push the span ceiling further. But all strategies work within the L^4 law — they improve stiffness efficiency, not the mode of stiffness generation.

4.4 The Ceiling: Self-Weight Versus EI

We can now answer precisely why girder bridges have a span limit.

When span increases at constant span-to-depth ratio, section area grows roughly as h^2 and self-weight as L^2. Moment of inertia grows as L^4. At first glance self-weight grows more slowly than stiffness. But self-weight bending moment grows as $L^2 \times L^2/8 = L^4/8$ — the fourth power. Stiffness growth (L^4) and moment growth (L^4) cancel three powers of L, leaving one:

> **Even at constant span-to-depth ratio, the self-weight deflection-to-span ratio still grows in proportion to span L.**

This “remaining L” is the fatal weakness of girder bridges. Every doubling of span doubles the self-weight deflection-to-span ratio. When it reaches the permissible value (e.g. L/500 or L/600), all stiffness has been consumed by self-weight and nothing remains for live load. This is the mechanical root of the girder bridge span limit: the competition between self-weight and EI — an irreconcilable conflict between the L^4 law and the growth rate of self-weight.

In practice: ordinary reinforced concrete simply-supported beams reach approximately 40 m. Prestressed concrete continuous box girders (all three strategies combined) reach 200–330 m, with the world record around 330 m. Steel box girders, with their higher strength-to-weight ratio, push somewhat further, but the limit remains. All these numbers are the direct consequence of the L^4-versus-self-weight competition.

4.5 The Strength of the Girder: Simplicity

Having discussed the limitations of girder bridges at length, the advantages deserve equal attention. The girder bridge is the most widely used bridge system: more than 90 per cent of the world’s bridges are girder bridges, and this is not without reason.

> **First, direct force flow.** The force path of a beam is the simplest of all systems: vertical load is transmitted to supports directly through bending. No intermediate members (cables, stays, towers, anchorages) are needed; no force conversion (from vertical to axial) is required. This directness means fewer members, simpler connections, and lower complexity.

Second, construction simplicity. Beams can be precast, cast in situ, cantilevered, or incrementally launched. Each method is inherently stable at every intermediate stage — a fundamental property stemming from the fact that a beam's stiffness does not depend on structural completeness. Each cantilevered segment, once cast, stands on its own; it does not need to wait for the entire bridge to close before becoming a structural system. An arch, by contrast, is unstable until closure; a suspension bridge's stiffening girder is incomplete until the cable is in place. This "local self-stability" of the beam is the mechanical root of its construction simplicity.

Third, predictable behaviour. The mechanical behaviour of a beam is the most predictable of all systems. No geometric nonlinearity (unlike cables), no compression instability (unlike arches), no multi-member coupling complexity (unlike cable-stayed systems). Engineers can have high confidence in a beam's behaviour — in many scenarios, this is extremely valuable.

These three advantages are the direct consequence of the pure-EI logic. Because stiffness comes entirely from the section, no additional members are needed. Because each segment independently provides stiffness, local self-stability is assured. Because geometric stiffness plays no part, behaviour is predictable. The advantages and limitations of the girder bridge share a single origin — yet another illustration that "a structural system is a bundled transaction".

4.6 The Ultimate Comparison of Four Systems

We can now re-examine the other systems from the girder bridge's vantage point, seeing exactly what each has broken through.

Arch bridge (Chapter 1): discovered that "shape can create stiffness" — converting bending into axial compression through curvature, replacing the section's bending capacity with its compressive capacity. Deflection control shifts from L^4/EI to L^2/EA — a two-order reduction in span exponent. The price: negative geometric stiffness, with greater load driving closer to instability.

Suspension bridge (Chapter 2): discovered that "geometry can create stiffness" — cable geometric stiffness bypasses the EI bottleneck entirely. Section parameters vanish from the deflection formula; stiffness is determined solely by geometric configuration and dead load. Self-weight transforms from pure burden to source of stiffness; the stiffening girder transforms from load-carrying member to enabling member.

Cable-stayed bridge (Chapter 3): discovered that "stiffness can be actively configured" — stays provide elastic supports with multiple independently adjustable parameters. The girder bridge's stiffness is "locked" by its section; the cable-stayed bridge's stiffness is "unlocked". Not a replacement of the stiffness source, but an increase in its controllability.

Feature	Girder	Arch	Suspension	Cable-Stayed
Stiffness source	Pure section EI	EA dominant + EI auxiliary	Positive K_G	Elastic supports + girder EI
Deflection dependence	L^4/EI	L^2/EA	Independent of section	$p/\bar{k}$ (stays govern)
Core parameter	Depth-to-span L/h	Rise-to-span f/L	Sag-to-span α	Spacing, tower ht., EA, force
Geometric stiffness	K_G = 0	K_G < 0 (harmful)	K_G > 0 (beneficial)	Mixed
Role of self-weight	Pure burden	Increases axial compression	Source of stiffness	Burden (compensated)
Role of EI	All stiffness	Stability guard	Enabling member	Active (high-order modes)
Adjustability	Very low	Low	Low	High
Advance over beam	— (baseline)	Shape creates stiffness	Geometry creates stiffness	Stiffness actively configurable
Price paid	Span limited	K_G < 0 → instability risk	Narrow adjustment range	System complexity

This table uses a unified stiffness language to bring four seemingly disparate bridge types into a single framework. Their essential distinction lies not in appearance but in the mechanism and mode of stiffness generation and configuration. From girder to arch to suspension to cable-stayed is the progressive deepening of humanity's understanding of how to configure K_E and K_G.

4.7 What the Beam Teaches Us

The stiffness logic of the girder bridge can be captured in one sentence: what you get depends entirely on what material you put in. No cable to amplify geometric stiffness, no stay to redistribute support stiffness, no curvature to convert bending into compression. You have only the section and its EI.

This purity is both limitation and strength. The limitation: the L^4 law and the competition with self-weight set an insurmountable span ceiling. The strength: simplicity — direct force flow, construction ease, predictable behaviour — makes it the optimal choice in the vast majority of scenarios.

From the perspective of Part One, the girder bridge provides the origin for understanding the entire stiffness spectrum. The geometric reconstitution of the arch, the geometric stiffness of the suspension bridge, the active configuration of the cable-stayed bridge — all are additional stiffness mechanisms superimposed on the beam's EI baseline. Understanding why the beam has limitations — why the

L^4 law cannot be broken, why section EI is “locked”, why self-weight must ultimately prevail — is the key to truly understanding what the other systems have innovated.

Without understanding the limitations of the beam, one cannot understand the price of the arch, the elegance of the suspension bridge, or the freedom of the cable-stayed bridge. The beam is the simplest of systems, but understanding it is the starting point for understanding all systems.

Part One concludes here. Part Two (Chapters 5–7) examines composite systems, boundary condition effects, and what makes a structure a system — all within the same stiffness framework.