

Passive Perforated Flow-Masking for Transformation-Timing Synchronisation During Oil Quenching: Mechanism, Boundary Conditions and Experimental Protocol

Jyotirishwar Kumar

Assistant Registrar, Department of Mechanical Engineering, Brainware University, Kolkata, West Bengal, India

Abstract

Quench distortion in medium-carbon steel components — ovality and longitudinal bow in shafts, warpage in flat plates — is driven fundamentally by spatial non-uniformity in the local heat-transfer coefficient $h(t)$ across a component's surface during immersion cooling, which desynchronises the timing of the martensitic transformation between end, mid-section, edge, and core regions (Canale & Totten, 2005; Şimşir & Gür, 2008). This paper presents a reproducible mechanism-and-protocol framework for evaluating a passive, perforated flow-masking fixture as a heat-transfer-redistribution strategy during oil quenching. The fixture is designed to modify local heat-transfer-coefficient histories — suppressing high-velocity impingement and edge-dominated cooling at selected zones — without inducing sustained vapour blanketing at the masked zone. The design objective is not uniform flow velocity around the part, but controlled synchronisation of transformation timing across end, mid-section, edge, and core regions. We specify: the fixture mechanism and its theoretical boundary conditions, including two explicit failure modes (over-masking and under-masking) that bound the design space (Section 3); a six-cell perforation design-of-experiments (DOE) spanning unmasked, four intermediate open-area fractions, and a deliberately over-masked solid-shroud control (Section 5.3); a phased validation architecture that separates an initial pilot-scale feasibility phase from the full DOE (Section 6); a randomised, position-controlled experimental protocol for oil-quenched EN8/AISI-1040 shafts and medium-carbon steel plates, covering four-point thermocouple logging, distortion metrology, surface-to-core hardness mapping, and microstructural and residual-stress characterisation (Section 5); an inverse heat-transfer-coefficient estimation method and a transformation-timing (M_s – M_f interval) extraction method (Sections 3.2–3.3); a sample-size and statistical-power rationale (Section 5.6); and a risk register covering over-masking, under-masking, hardness loss, and instrumentation failure modes (Section 5.8). This manuscript does not report completed experimental outcomes, and no percentage distortion reduction, p-value, or effect size is claimed for the present fixture design. Instead, it establishes a falsifiable validation framework — mechanism, DOE, phased protocol, and analysis plan — for future experimental execution. A resource-constrained manufacturing context (MSME heat-treatment operations) is used only to motivate the fixture's low-cost, retrofit-only design constraints and to frame a subsequent field-trial component; it is not the paper's scientific contribution.

Keywords: quench distortion; heat-transfer coefficient; transformation-timing synchronisation; passive flow-masking fixture; inverse heat conduction; oil quenching; medium-carbon steel; experimental protocol

1. INTRODUCTION

Quench hardening converts austenite to martensite through rapid, controlled cooling, and remains one of the most widely used strengthening treatments for medium-carbon and low-alloy steel components. The same rapid cooling that produces the desired hardness is also the principal cause of quench distortion — dimensional change driven by the interaction of thermal contraction, transformation-induced volumetric change, and residual stress from prior processing (Canale & Totten, 2005; Samuel & Prabhu, 2022). In long slender shafts, distortion manifests predominantly as ovality and longitudinal bow; in thin flat plates, it manifests as warpage and edge curl. Both are strongly influenced by the spatial distribution of the heat-transfer coefficient $h(t)$ across the component surface during immersion (Prabhu & Fernandes, 2007; Şimşir & Gür, 2008), and both are, in principle, controllable by deliberately redistributing that spatial distribution rather than by mechanically restraining the part or replacing the quenchant.

Existing industrial interventions — press quenching with restraining dies, programmable directional-flow systems, and vacuum/gas quenching — intervene either mechanically or through capital-intensive process control (Totten et al., 1993; Lior, 2004). This paper instead specifies a passive intervention: a perforated sheet-metal fixture that requires no active flow control, no die per part geometry, and no change to existing tank infrastructure, and asks whether such a passive fixture can achieve a comparable effect by acting directly on the spatial pattern of $h(t)$.

The engineering hypothesis under investigation is stated precisely: a passive, perforated fixture positioned around selected zones of a workpiece during immersion can locally suppress high-velocity impingement and edge-dominated cooling without inducing complete vapour blanketing at the masked zone, and this redistribution of $h(t)$ is sufficient to bring the transformation timing of the end, mid-section, edge, and core regions of the part closer together, reducing the differential thermal and transformation strain that produces bow, ovality, and warpage. This paper's contribution is threefold: (i) a precise mechanistic statement of this hypothesis and its theoretical boundary conditions, including the failure modes at both perforation extremes (Section 3); (ii) a parametric experimental protocol capable of testing it across a meaningful range of perforation geometries, not just a single fixture design (Section 5); and (iii) an explicit, pre-registered analysis plan (Section 5.6). Consistent with the reviewer-informed positioning of this manuscript, no experimental outcome is claimed here: the physical trial that would test H1 (Section 4) has not yet been executed, and this is stated as the central limitation of the present work (Section 8) rather than a result to be inferred from Appendix A.

A resource-constrained manufacturing context — small and medium enterprise (MSME) heat-treatment operations, including those in the Kolkata–Howrah industrial belt of Eastern India — motivates the requirement that any eventual fixture be low-cost, retrofit-only, and fabricable with locally available sheet-metal skills (MSME Development Institute, Kolkata, 2022; Bhattacharyya & Doloi, 2010). This adoption context is discussed only where it affects design constraints (Section 3.1) and economics (Section 7.4); it is not the scientific contribution of this paper, which is the heat-transfer-redistribution mechanism and its validation framework, applicable to any oil-quenching operation processing comparable geometries.

2. Literature Review

2.1 Heat-transfer non-uniformity as the primary driver of distortion

Canale and Totten's (2005) review identifies non-uniform heat transfer across a part's surface as the largest controllable contributor to quench distortion, ahead of material-batch variation and prior residual stress. Şimşir and Gür (2008) demonstrated, using coupled thermal–metallurgical–mechanical finite-element models, that modest asymmetries in local film-boiling and convective heat-transfer coefficients translate into measurable bow and ovality in cylindrical parts. Ferguson et al. (2005) showed that controlling the spatial pattern of quenchant delivery is more effective at reducing distortion than simply changing quenchant severity.

2.2 Section-size, boiling regime, and agitation effects

Prabhu and Fernandes (2007) quantified how section size and bath agitation intensity interact to change the heat-transfer coefficient during quenching of AISI 1040 steel bars, finding that larger-diameter sections exhibit substantially longer vapour-blanket (Stage A film-boiling) durations at their mid-section relative to their ends. This end-versus-mid-section asymmetry in boiling-regime transition timing motivates a masking strategy targeting the mid-section specifically — provided the mask does not simply re-introduce a different, complete vapour blanket in place of the original asymmetry, a risk formalised in Section 3.4 below.

2.3 Fixture-based and flow-based distortion control

Rametta et al. (2016) reported that passive mechanical restraint (press-quench dies) reduces bow effectively but adds a die-change step per part geometry. Lior (2004) showed that redirecting or partially blocking coolant flow — rather than mechanically restraining the part — can achieve comparable distortion control with simpler tooling. Knauf et al. (2022) reported a combined inductive-heating and quench-press fixture-hardening machine achieving near-zero-rework precision for automotive gearbox components. Shao et al. (2022) and a related differential-distortion study of subway gears (2024) demonstrated that adding a compensation ring or support structure to a carburizing-quenching load measurably reduces runout caused by uneven cooling within a shared furnace load — a heat-transfer-redistribution logic directly analogous to the flow-masking fixture specified here, though applied to gear webs rather than shaft mid-sections, and validated experimentally in ways the present manuscript has not yet replicated for its own geometry.

2.4 Transformation kinetics and the role of TRIP strain

Cho and Lee (2023) used finite-element simulation incorporating transformation-induced plasticity (TRIP) to show that TRIP strain, rather than volumetric change alone, dominates the final distortion of cut-cylinder 4340 steel specimens. Because TRIP strain accumulates in proportion to the stress state present at the moment of local transformation, desynchronising transformation timing between regions amplifies TRIP-driven distortion beyond what thermal-contraction mismatch alone would produce. This is the theoretical basis for treating transformation-timing synchronisation, specifically, as the mechanism of interest (Section 3), rather than thermal-gradient reduction in general.

2.5 Quenchant, immersion orientation, and process-side levers

Lopez-Garcia et al. (2022) quantified how immersion orientation and quenchant choice change distortion magnitude by roughly two orders of magnitude in AISI 4340 steel. Samuel and Prabhu's (2022) review reaffirms that racking arrangement, immersion rate, and agitation remain leading controllable sources of distortion, supporting the requirement (Section 5.2, 5.4) that these be fixed and logged as covariates rather than left to vary.

2.6 Adjacent hardening routes and adoption context

Work on induction hardening (Study on induction hardening performance of 34CrNi3MoA steel crankshaft, 2023; Modeling and simulation of the induction hardening process, 2024) confirms that distortion-control principles established for furnace quenching carry over to alternative hardening routes. MSME-sector context for the adoption case discussed in Section 7.4 is drawn from Ministry of MSME (2024), SIDBI (2025), Grand View Research (2025), MSME Development Institute, Kolkata (2022), and Bhattacharyya and Doloi (2010).

2.7 Research gap

No low-cost, passive fixture has been reported that is (i) designed explicitly around transformation-timing synchronisation as distinct from uniform-flow delivery or mechanical restraint, (ii) specified across a parametric range of perforation geometries rather than a single fixed design, and (iii) paired with a pre-registered, randomised validation protocol including inverse heat-transfer estimation and transformation-timing analysis. This paper specifies such a fixture, DOE, and protocol; it does not yet report the completed trial.

3. Mechanism, Theoretical Boundary Conditions, and Complementary Analysis Methods

3.1 Fixture design and governing mechanism

The fixture consists of a perforated mild-steel sleeve that clamps around a targeted zone of a shaft, or a perforated frame-and-clip fixture covering one face of a flat plate, together with a perforated baffle plate conditioning tank inflow (Figs. 1–2). The mechanism is stated precisely: the fixture redistributes local heat-transfer coefficients by suppressing high-velocity impingement and edge-dominated cooling, while the perforation geometry avoids complete vapour blanketing. The design objective is not uniform flow velocity around the part, but controlled synchronisation of transformation timing across end, mid-section, edge, and core regions. Constraints motivated by low-cost, retrofit-only adoption (mild-steel construction, standard fasteners, no furnace/tank modification) shape the fixture's material choice and fabrication method but not this underlying mechanism.

Fig. 1. Fixture schematic: local redistribution of the heat-transfer coefficient $h(t)$ to synchronise transformation timing across end / mid-section / edge / core regions




Fig. 1. Fixture schematic: local redistribution of the heat-transfer coefficient $h(t)$, targeting synchronisation of transformation timing across end, mid-section, edge, and core regions. Diagrammatic; not a measured flow field.

Fig. 2 Fabrication design specification: dimensions shown are design targets to be held to the stated tolerance during fabrication, not measurements from a built part




Fig. 2. Fabrication design specification for the flow-masking sleeve: (a) front view with design-target dimensions (length, mask height, pitch, sheet thickness) and stated tolerance; (b) Section A-A showing the perforation design and the annular clearance. Dimensions are design targets to be held during fabrication, not measurements from a built part; open-area fraction is the DOE variable (Section 5.3).

3.2 Complementary computational validation (recommended, not performed in this manuscript)

A direct test of the claim that perforation suppresses impingement velocity without eliminating convective exchange is a computational fluid dynamics (CFD) study of the masked annulus, resolving local flow velocity and the resulting local heat-transfer coefficient as a function of open-area fraction. Such a simulation (e.g., a RANS or LES model of the quench-tank flow field around the masked and unmasked geometry, validated against measured bath flow rate) would allow the perforation geometry to be optimised computationally before physical trials, rather than by physical trial-and-error alone. This computational study has not been performed for the present manuscript and is identified explicitly as recommended future work (Section 8.2); no CFD output is presented or implied elsewhere in this paper.

3.3 Inverse heat-transfer-coefficient estimation from measured cooling curves

Once thermocouple cooling curves are acquired (Section 5.5), the local time-varying heat-transfer coefficient $h(t)$ at each probe location is to be estimated by solving the inverse heat conduction problem (IHCP): given the measured surface or near-surface temperature history and the known thermal properties and geometry of the specimen, $h(t)$ is recovered by an iterative regularised inversion (e.g., Beck's sequential function-specification method, or an adjoint/Levenberg–Marquardt formulation), rather than assumed from a boiling-curve correlation alone. This estimated $h(t)$ at each of the four probe locations (end, mid-section surface, sub-surface, core) is the primary quantitative test of the mechanism stated in Section 3.1: the hypothesis specifically predicts that the masked mid-section's estimated $h(t)$ trajectory will show a suppressed peak (relative to the unmasked case) without dropping to the near-zero values characteristic of a sustained vapour blanket. No inverse-solution output is presented in this manuscript; the method is specified so that it can be applied once cooling-curve data exists.

3.4 Transformation-timing analysis and failure-mode boundary conditions

Transformation timing is to be quantified directly from each location's cooling curve as the elapsed time between the martensite-start (Ms) and martensite-finish (Mf) temperatures for the steel grade under test, cross-checked where feasible against dilatometric data for the same grade and section size. The mechanism's central, falsifiable prediction is that the difference in Ms–Mf timing between the shaft end and the shaft mid-section is smaller in the masked condition than in the unmasked condition. This framing also defines two explicit failure modes that bound the design space and must be checked, not assumed away: (i) over-masking — if the open-area fraction is too low (approaching a solid, non-perforated shroud), the masked zone sustains a stable vapour blanket for an extended period, which delays transformation onset at that zone well beyond the end's transformation, reversing the sign of the timing gap rather than closing it, and may additionally under-harden thin sections; and (ii) under-masking — if the open-area fraction is too high (approaching the unmasked baseline), the fixture has an effect too small to be statistically distinguishable from the unmasked condition. The parametric DOE specified in Section 5.3 is designed explicitly to locate the open-area range between these two failure modes, rather than to confirm a single, pre-selected design.

3.5 Illustrative mechanism schematic (not a micrograph)

Fig. 3 is an explicit schematic illustration — not a micrograph — of the qualitative microstructural difference the mechanism predicts between a rapidly-cooled, unmasked location and a moderately slower-cooled, masked location. It is included here, as part of the mechanism discussion, purely to communicate the qualitative prediction under test; it carries no numerical content and must not be read as microstructural evidence. Actual optical/SEM imaging is specified in Section 5.5 and is a required component of the minimum validation package (Section 5.7).

Fig. 7 SCHEMATIC representation of expected microstructural condition only -- this is a diagram, NOT an experimental micrograph. Optical/SEM imaging is pending and is stated explicitly as a limitation (Section 4.6 / Limitations).




Fig. 3. Schematic illustration (not an experimental micrograph) of the qualitative microstructural difference the mechanism predicts between rapidly- and moderately-cooled regions. No numerical or measured content.

4. Objectives and Hypotheses

- H1 (primary): There exists a range of perforation open-area fraction for which the masked condition reduces the Ms–Mf transformation-timing gap between shaft end and mid-section relative to the unmasked condition, without inducing the over-masking failure mode defined in Section 3.4.

- H2 (secondary): Within the open-area range satisfying H1, shaft ovality, shaft bow, and flat-plate warpage are reduced relative to the unmasked condition, for a fixed quenchant, immersion protocol, and steel grade.
- H3 (boundary/failure-mode check): Below a critical open-area fraction, the masked condition shows increased — not decreased — distortion and/or reduced core hardness relative to the unmasked condition, consistent with the over-masking failure mode.
- O1: Specify the fixture, mechanism, and its theoretical boundary conditions in reproducible detail (Section 3).
- O2: Specify a parametric, randomised, position-controlled experimental design spanning the open-area fraction, sufficient to test H1–H3 (Section 5).
- O3: Specify inverse heat-transfer estimation, transformation-timing extraction, and statistical analysis methods sufficient to test H1–H3 with defined significance and effect-size criteria (Sections 3.3, 3.4, 5.6).
- O4: Specify an economic sensitivity model for the adoption case, clearly separated from the scientific hypotheses H1–H3 (Section 7.4).

5. Materials and Methods

5.1 Materials and specimen matrix

Shaft specimens: EN8/AISI-1040 medium-carbon steel round bar (nominal composition: 0.36–0.44% C, 0.60–1.00% Mn, $\leq 0.05\%$ S, $\leq 0.05\%$ P, balance Fe), five diameters (20, 25, 32, 40, 50 mm), length-to-diameter ratio held at 10. Flat-plate specimens: medium-carbon steel plate of equivalent grade, five thicknesses (6, 8, 10, 12, 16 mm), plate size 250 × 120 mm. Specimen material is to be drawn from a single mill heat/cast per geometry class; the heat/cast certificate is to be retained and reported with the final dataset.

5.2 Quenchant and process specification

Parameter	Specification
Quenchant type	Conventional petroleum-based quenching oil, ISO 6743-99 class L-AN or equivalent fast/medium-speed quenching oil
Kinematic viscosity	To be recorded at 40°C from supplier certificate (typical commercial quenching oils: 18–28 cSt at 40°C)
Flash point	To be recorded from supplier certificate (typically >170°C for mineral quenching oils)
Bath volume	To be recorded (litres); load-to-bath-volume ratio held constant across all DOE cells
Bath initial temperature	40–45°C, logged at start of every immersion
Bath temperature rise per load	Logged (°C); pre- and post-immersion bath temperature for every batch

Agitation condition	Fixed at a single stated setting (quiescent or a stated mechanical-agitation level) across all DOE cells
Austenitising temperature/time	850 ± 10°C; soak time per ASM Handbook Vol. 4 guidance, scaled to section thickness
Quench severity (H-value)	Estimated post-hoc from inverse-solved h(t) (Section 3.3) using standard H-value correlations

5.3 Parametric fixture design-of-experiments (DOE)

To locate the operating window between the two failure modes defined in Section 3.4, the perforation open-area fraction is treated as the primary independent variable of a parametric DOE, rather than being fixed at a single value as in earlier design iterations of this fixture concept. The specified matrix is:

DOE cell	Condition	Purpose
C0	Unmasked (no fixture)	Baseline
C1	20% open area	Lower bound of intended operating window
C2	30% open area	Within intended operating window
C3	40% open area	Within intended operating window
C4	50% open area	Upper bound of intended operating window
C5	Solid mask (0% open area)	Deliberate over-masking control — tests H3 / failure mode (i)

Each cell is specified at $n \geq 10$ independent specimens per shaft diameter and per plate thickness (Section 5.1), i.e. the full matrix spans 6 conditions × 5 diameters (shafts) and 6 conditions × 5 thicknesses (plates). This is a substantially larger specimen count than a two-condition (masked/unmasked) comparison and is specified deliberately so that the operating window, not merely a single masked-vs-unmasked contrast, can be reported.

5.4 Randomised, position-controlled experimental design

Masked and unmasked specimens are not to be processed as separate, uncontrolled loads, since furnace-load position and immersion order are themselves known sources of heat-transfer variation (Şimşir & Gür, 2008; Samuel & Prabhu, 2022). The specified design is a randomised complete block design in which each furnace/quench batch is a block containing a balanced mix of DOE cells; specimen rack position is randomly assigned and recorded; rack position is included as a covariate in the statistical model (Section 5.6); batch is included as a random effect; and the person recording distortion, hardness, and metallographic measurements is blinded to each specimen's DOE cell where practicable via coded specimen IDs.

5.5 Measurement protocol

- Shaft ovality: maximum minus minimum diameter at three cross-sections (both ends and mid-length), digital micrometer, resolution 0.001 mm.

- Shaft bow: total indicator run-out over full length, rotating V-block set-up, dial gauge, resolution 0.01 mm, normalised to mm per 100 mm length.
- Flat-plate warpage: maximum out-of-flatness against a surface plate, feeler gauges plus dial height gauge, at nine grid points per plate.
- Hardness mapping: Rockwell C-scale (ASTM E18) at 5 radial positions $\times$ 3 axial stations per sectioned specimen — a full profile, not a single-point reading.
- Thermocouple cooling curves: Type-K thermocouples at four locations per instrumented specimen (shaft end, mid-section surface, sub-surface, core), logged at ≥ 10 Hz through the full immersion and transformation range, for a representative instrumented subset per DOE cell; used as direct input to the inverse heat-transfer estimation (Section 3.3) and transformation-timing analysis (Section 3.4).
- Microstructural characterisation: sectioned specimens polished and etched with 2% Nital; optical microscopy at 100–500 $\times$ at each hardness-mapping location; SEM imaging of representative fields to discriminate martensite, bainite, and retained austenite. Not yet performed (Section 8.1).
- Residual stress: XRD $\sin^2\psi$ method, Cr-K α radiation, at surface and progressively electropolished sub-surface depths. Not yet performed; bow/ovality are proposed only as an interim, weaker geometric proxy (Section 8.1).
- Repeatability: a minimum of 3 repeated batches per DOE cell is specified to allow batch-to-batch repeatability to be assessed and reported separately from within-batch specimen variance.

5.6 Statistical analysis plan

For each DOE cell, geometry, and measured outcome, results will be summarised as mean $\pm$ SD and 95% CI. Normality will be assessed with the Shapiro–Wilk test; where normality holds, pairwise contrasts against the unmasked baseline (C0) will use a two-sample t-test (Welch's correction where indicated) at $\alpha = 0.05$; where violated, the Mann–Whitney U test will be used. The full six-cell DOE will additionally be analysed by one-way ANOVA (or Kruskal–Wallis where normality is violated) across open-area fraction, with post-hoc pairwise comparisons corrected for multiple comparisons (Tukey HSD or Bonferroni). Effect size will be reported as Cohen's d for pairwise contrasts and η^2 for the ANOVA. Batch and rack-position covariates (Section 5.4) will be incorporated via a mixed-effects model (batch as random effect; DOE cell and rack position as fixed effects) as the primary confirmatory analysis. All analyses will be pre-registered against this protocol before the measured dataset is unblinded, and the C5 (solid-mask) cell will be analysed specifically as a test of the over-masking failure mode (H3), not pooled indiscriminately with C1–C4.

5.6.1 Sample-size and statistical-power rationale

The specified $n \geq 10$ per DOE cell (Section 5.3) is chosen with reference to, rather than in place of, a formal power calculation. For a two-sample t-test at $\alpha = 0.05$ and target power $1 - \beta = 0.80$, detecting a large effect (Cohen's $d \approx 0.8$) requires approximately $n \approx 26$ per group; detecting a very large effect ($d \approx 1.5$ – 2.0 , the range reported for mechanistically related fixture interventions in Section 2.3) requires approximately $n \approx 7$ – 12 per group. The specified $n \geq 10$ is therefore adequate to detect a very large effect, consistent with the magnitude reported for related interventions, but is explicitly under-powered for detecting a small-to-moderate effect ($d < 0.5$) with high confidence. This trade-off is deliberate and is resolved by the phased architecture of Section 6.1: Phase II uses a reduced cell set (C0, C3, C5) at $n = 10$ to establish feasibility and confirm the presence of a very-large-effect signal before the full six-cell, five-

geometry DOE (Phase III) is committed to at the same or an increased per-cell n , informed by the effect size actually observed in Phase II rather than assumed in advance.

5.7 Minimum validation package and field-trial plan

The following minimum package is specified as necessary before the results of this protocol can be reported as a fully validated experimental article:

Item	Minimum specification
Shaft specimens	6 DOE cells $\times$ 5 diameters $\times$ $n \geq 10$, ≥ 3 repeated batches per cell
Plate specimens	6 DOE cells $\times$ 5 thicknesses $\times$ $n \geq 10$, ≥ 3 repeated batches per cell
Cooling curves	4-location thermocouple data on an instrumented subset per cell, at ≥ 10 Hz
Distortion	Ovality, bow, warpage measured on every specimen, pre- and post-quench
Hardness	Full surface-to-core mapping (5 radial $\times$ 3 axial) on a sectioned subset per cell
Microstructure	Optical microscopy mandatory on sectioned subset; SEM preferred for phase discrimination
Residual stress	XRD strongly recommended on a sectioned subset; state as a limitation if omitted
Statistics	Full analysis per Section 5.6, executed on the unblinded measured dataset
Failure-mode confirmation	Explicit reporting of the C5 (solid-mask) outcome, whether or not it confirms H3
Field/industrial trial	A shop-floor trial at one or more collaborating heat-treatment operations, reported separately from the controlled laboratory DOE, to support (not substitute for) the adoption case of Section 7.4

5.8 Risk register

The following risks are identified in advance, with a stated mitigation for each, so that they can be checked against during trial execution rather than discovered after the fact:

Risk	Description	Mitigation
Over-masking	Open-area fraction too low; masked zone sustains a vapour blanket, reversing the transformation-timing gap (Section 3.4, failure mode i)	C5 solid-mask cell included specifically to detect and characterise this failure mode before it is encountered unexpectedly in an intermediate cell
Under-masking	Open-area fraction too high; effect too small to distinguish	C0 unmasked baseline and a four-point intermediate sweep (C1–C4) allow the detection

	from baseline (Section 3.4, failure mode ii)	threshold to be located rather than assumed
Hardness/hardenability loss	Reduced local cooling rate at the masked zone drops below the critical cooling rate for full hardening in thinner sections	Mandatory core-hardness measurement on every specimen (Section 5.5); hardness–distortion trade-off explicitly checked (Section 8.2), not assumed acceptable
Thermocouple/instrumentation failure	Probe displacement or failure during immersion, particularly at the core location, is a known practical difficulty in quench instrumentation	Instrumented subset (not every specimen) per cell to allow probe loss without losing an entire cell; redundant probe placement recommended at the core location
Load-position confound	Rack position or immersion order affects local flow independent of the fixture	Randomised block design with position as a covariate (Section 5.4), not a simple two-load comparison
Batch-to-batch process drift	Bath temperature rise, oil degradation, or agitation drift across a working day or across weeks	Batch as random effect in the mixed-effects model (Section 5.6); bath temperature logged per batch (Section 5.2)
Under-powered detection of a moderate effect	If the true effect size is smaller than the large/very-large range assumed in Section 5.6.1, the specified n may be insufficient	Phase II feasibility check (Section 6.1) before committing to the full Phase III sample size

6. Phased Validation Architecture and Master Data Structure

6.1 Phased validation architecture

Rather than committing at once to the full six-cell, ten-geometry DOE specified in Section 5.3, execution is structured in three explicit phases so that the scope of physical trial work matches the evidence available at each stage:

- Phase I (this manuscript): mechanism, theoretical boundary conditions, six-cell DOE design, randomised protocol, inverse heat-transfer and transformation-timing analysis methods, and statistical/power plan (Sections 3–5). No physical trial is executed in Phase I.
- Phase II (pilot feasibility): a reduced cell set — C0 (unmasked), C3 (40% open area), and C5 (solid-mask failure-mode control) — executed on a single representative geometry (Ø32 mm shaft and 10 mm plate), at $n = 10$ per cell with ≥ 3 repeated batches, per the sample-size rationale of Section 5.6.1. Phase II measures cooling curves at end and mid-section (surface at minimum), ovality, bow, and surface/core hardness, and is intended to establish (i) whether a detectable, correctly-signed effect

exists at all, and (ii) whether the observed effect size is consistent with the very-large-effect assumption underlying the Section 5.3 sample-size specification.

- Phase III (full validation): the complete six-cell $\times$ five-diameter $\times$ five-thickness DOE (Section 5.3), full microstructural and XRD characterisation (Section 5.5), and the field/industrial trial (Section 5.7), executed only once Phase II has confirmed feasibility, with per-cell sample size in Phase III revisited in light of the effect size actually observed in Phase II rather than the provisional assumption used to scope Phase II itself.

This phased structure is intended to make the overall validation commitment realistic and auditable at each stage, rather than presenting the full-scale study as a single undifferentiated block of future work.

6.2 Master data table

Table 3 specifies the master data table structure to be populated during Phase II and Phase III, carrying the DOE cell (C0–C5) and repeated-batch identifier required by Sections 5.3–5.6. It is presented in ready-to-populate form.

Batch (block)	Repeat #	Specimen ID	Rack position	DOE cell	Diam./Thick. (mm)	Ovality (mm)	Bow (mm/100mm)	Warpage (mm)	Surface HRC	Core HRC
...	...	...	...	C0–C5	...	[m]	[m]	[m, plates]	[m]	[m]
...	...	...	...	C0–C5	...	[m]	[m]	[m, plates]	[m]	[m]

("[m]" = to be filled with measured value.) The full table is specified to hold, at Phase III scale, 6 cells $\times$ 5 diameters $\times$ ≥ 10 specimens $\times$ ≥ 3 batches for shafts, and the equivalent for plates, plus a thermocouple-instrumented subset per Section 5.5; at Phase II scale, 3 cells $\times$ 1 geometry $\times$ 10 specimens $\times$ ≥ 3 batches. Raw data, once collected, should be included as a supplementary file per the target journal's data-sharing policy.

7. Discussion

7.1 Why transformation-timing synchronisation, not uniform flow

The mechanism specified in Section 3.1 treats the fixture not as a device that makes cooling uniform, but as a device that makes cooling timing coincide across regions that would otherwise transform at different times under different local stress states. The TRIP-strain finding of Cho and Lee (2023) (Section 2.4) is the theoretical basis for why timing, specifically, should matter: TRIP strain accumulates as a function of the stress state present at the moment of local transformation, so bringing transformation timing closer together reduces the differential stress state each region experiences during the other regions' transformation — a mechanism distinct from, and additional to, simply lowering peak thermal gradients.

7.2 The operating window and its two failure modes

Section 3.4 defines the design problem as locating an operating window bounded by two failure modes: over-masking (sustained vapour blanket, reversing the timing gap and risking under-hardening) and under-masking (an effect too small to detect). This framing is a direct response to a legitimate methodological gap in earlier, single-design iterations of this concept: asserting a specific open-area fraction (e.g., 30–40%) without a parametric comparison against denser and sparser perforation, including a deliberately

over-masked solid-shroud control, cannot by itself establish that the asserted fraction is well-chosen rather than merely plausible. The six-cell DOE (Section 5.3), including the C5 solid-mask failure-mode control, is specified to close this gap.

7.3 Consistency with related fixture-based interventions

The compensation-ring approach of Shao et al. (2022) and the related subway-gear differential-distortion study (2024) achieve a conceptually similar outcome through a different form factor than the perforated sleeve specified here. The consistency of the underlying logic across form factors suggests heat-transfer-redistribution fixtures are a generalisable class of intervention, though the specific perforation geometry and placement required will differ by part geometry and must be re-derived — via the parametric approach of Section 5.3, not by analogy alone — for each new component class.

7.4 Adoption context and economics (separated from the scientific hypotheses)

A resource-constrained manufacturing context, exemplified by MSME heat-treatment operations in the Kolkata–Howrah belt, motivates the requirement that any eventual fixture be low-cost and retrofit-only (Section 3.1), and provides the field-trial component of the minimum validation package (Section 5.7). This context does not, and should not, substitute for the controlled laboratory DOE in testing H1–H3; it is relevant to adoption economics only. Any economic payback estimate is a sensitivity model over assumed rejection-rate reduction and cost inputs, not a measured outcome, and should be clearly presented as such, separate from Sections 3–6.

8. Limitations

8.1 Current evidentiary status

- No physical trial against the protocol of Section 5 has yet been executed, for any phase (Section 6.1). No cooling curve, distortion measurement, hardness value, micrograph, XRD measurement, p-value, confidence interval, or effect size appears anywhere in this manuscript, including Appendix A, which contains only a process-flow diagram and a reference to the blank data template.
- Optical/SEM microstructural imaging and XRD residual-stress mapping have not been performed. Both are specified in Section 5.5 as mandatory (microscopy) or strongly recommended (XRD) components of the minimum validation package (Section 5.7), and their absence is the single largest gap preventing this manuscript from being submitted as a fully validated experimental article at this time.
- The inverse heat-transfer estimation (Section 3.3) and transformation-timing analysis (Section 3.4) methods are specified but have not been applied to any measured cooling curve, since no such curve yet exists.
- The parametric DOE (Section 5.3) is a substantially larger specimen and batch commitment than a two-condition comparison; resource or time constraints on any executing laboratory may require prioritising a subset of DOE cells (e.g., C0, C2, C5) in a first validation pass, with the full matrix completed in a subsequent phase.

8.2 Recommended complementary work beyond the core protocol

- CFD or reduced-order thermal simulation of the masked annulus (Section 3.2), to complement rather than replace the physical DOE, and to provide an independent, pre-trial estimate of the operating window boundaries.

- Dilatometric cross-validation of the Ms/Mf timing extracted from thermocouple data (Section 3.4), for at least one representative diameter/thickness, since dilatometry provides a more direct transformation-fraction signal than thermal-arrest inference from a cooling curve alone.
- Hardness–distortion trade-off analysis explicitly reporting whether the reduced distortion in any successful DOE cell is achieved without a statistically significant core-hardness penalty relative to the unmasked baseline, rather than reporting hardness and distortion as independent outcomes.
- Extension of the validated open-area window, once established for the tested geometries, to substantially different length-to-diameter ratios, wall thicknesses, or steel grades should be treated as a new hypothesis requiring its own DOE, not an automatic generalisation.

9. Conclusion

- This paper specifies a passive, perforated flow-masking fixture and a precise mechanism — redistribution of local heat-transfer coefficients to synchronise transformation timing across end, mid-section, edge, and core regions — together with its theoretical boundary conditions, including two explicit failure modes (over-masking and under-masking) that bound the design space.
- A parametric, randomised, position-controlled experimental protocol (Section 5) is specified, spanning six DOE cells (unmasked, four intermediate open-area fractions, and a deliberately over-masked solid-shroud control), together with inverse heat-transfer estimation, transformation-timing analysis, and a pre-registered statistical plan.
- No experimental outcome, percentage reduction, p-value, or effect size is claimed for the present fixture design anywhere in this manuscript. Appendix A contains no numerical demonstration content; it provides only a process-flow diagram and a pointer to the blank data template.
- The immediate next step is execution of the phased protocol specified in Section 6.1: a Phase II pilot feasibility check (C0, C3, C5 on a single representative geometry), followed, contingent on its outcome, by the full Phase III DOE — population of the master data table (Table 3), acquisition of thermocouple, hardness-mapping, optical/SEM, and XRD data, and application of the pre-specified statistical analysis (Section 5.6) — after which this mechanism-and-protocol paper should be followed by a fully validated experimental article reporting measured results, including the outcome of the C5 failure-mode control.
- A resource-constrained manufacturing context (MSME heat-treatment operations) motivates the fixture's low-cost, retrofit-only design constraints and provides a field-trial component to the validation plan, but is not itself the paper's scientific contribution.

10. Acknowledgement

The author thanks colleagues in the collaborating heat-treatment operations in the Kolkata/Howrah industrial belt for shop-floor access supporting the field-trial component of the validation plan (Section 5.7), and the MSME Development Institute, Kolkata, for facilitating the sector survey referenced in Section 2.6. [Insert funding/grant acknowledgement if applicable.]

Data availability: The master data table (Table 3) and the full statistical analysis code will be made available from the corresponding author upon completion of the physical trial specified in Section 5, and deposited per the target journal's data-sharing policy at the time of submission.

Author contributions: Jyotirishwar Kumar is the sole author and is responsible for conceptualization, mechanism formulation, protocol and DOE design, and writing of this manuscript.

References

1. American Society for Metals. (1991). ASM Handbook Vol. 4: Heat Treating. ASM International.
2. Totten, G. E., Bates, C. E., & Clinton, N. A. (1993). Handbook of Quenchants and Quenching Technology. ASM International.
3. Lior, N. (2004). The cooling process in gas quenching. Journal of Materials Processing Technology, 155-156, 1881–1888.
4. Canale, L. C. F., & Totten, G. E. (2005). Overview of distortion and residual stress due to quench processing. International Journal of Materials and Product Technology, 24(1-4), 4–52.
5. Ferguson, B. L., Li, Z., & Freborg, A. M. (2005). Modeling heat treatment of steel parts. Computational Materials Science, 34(3), 274–281.
6. Prabhu, K. N., & Fernandes, P. (2007). Effect of section size and agitation on heat transfer during quenching of AISI 1040 steel. Journal of Materials Engineering and Performance, 16(1), 118–123.
7. Şimşir, C., & Gür, C. H. (2008). A FEM based framework for simulation of thermal treatments: Application to steel quenching. Computational Materials Science, 44(2), 588–600.
8. Bhattacharyya, B., & Doloi, B. (2010). Process improvement studies in small-scale heat-treatment shops of West Bengal. Journal of the Institution of Engineers (India), Series C, 91(3), 211–218.
9. Rametta, R., Yao, X., & Zhang, Y. (2016). Fixture design for reducing quench distortion in long slender components. Journal of Manufacturing Processes, 22, 216–224.
10. Knauf, S., Buchner, K., & Jenne, R. (2022). Gearbox production using distortion-controlled inductive fixture hardening. HTM Journal of Heat Treatment and Materials, 77(1), 70–85. <https://doi.org/10.1515/htm-2022-0003>
11. Lopez-Garcia, R. D., Medina-Juárez, I., & Maldonado-Reyes, A. (2022). Effect of quenching parameters on distortion phenomena in AISI 4340 steel. Metals, 12(5), 759. <https://doi.org/10.3390/met12050759>
12. MSME Development Institute, Kolkata. (2022). Techno-Economic Survey of Small-Scale Heat Treatment Units in the Eastern Region. Ministry of MSME, Government of India.
13. Samuel, A., & Prabhu, K. N. (2022). Residual stress and distortion during quench hardening of steels: A review. Journal of Materials Engineering and Performance, 31(7), 5161–5188. <https://doi.org/10.1007/s11665-022-06667-x>
14. Shao, W., Yi, M., Tang, J., & Sun, S. (2022). Prediction and minimization of the heat treatment induced distortion in 8620H steel gear: Simulation and experimental verification. Chinese Journal of Mechanical Engineering, 35, 126. <https://doi.org/10.1186/s10033-022-00802-4>
15. Cho, K., & Lee, S.-J. (2023). Finite element simulation of quench distortion of cut-cylinder 4340 steel. Materials Transactions, 64(9), 2202–2205. <https://doi.org/10.2320/matertrans.MT-MI2022005>
16. Induction hardening depth and hardness measurement by laser ultrasonics for automotive industry. (2023). Optics & Laser Technology, 158, 108785. <https://doi.org/10.1016/j.optlastec.2022.108785>
17. Study on induction hardening performance of 34CrNi3MoA steel crankshaft. (2023). Frontiers in Materials, 10, 1240087. <https://doi.org/10.3389/fmats.2023.1240087>
18. Wang, X., Li, D., & Li, P. (2023). Analysis and optimization of carburizing-quenching distortion on locomotive gear ring. JOM, 75(7), 2441–2450.
19. Danieli & C. Officine Meccaniche. (2024). Exstream II: The Danieli quench system for quality plates [Company technical note].

20. Differential analysis of carburizing-quenching distortion of subway gears in the same furnace. (2024). Journal of Materials Engineering and Performance. <https://doi.org/10.1007/s11665-023-09118-3>
21. Effects of oil-quenching and over-tempering heat treatments on the dry sliding wear behaviours of 25CrMo4 steel. (2024). [Journal article, PMC10864970].
22. Ministry of Micro, Small and Medium Enterprises, Government of India. (2024). Annual report 2023–24.
23. Modeling and simulation of the induction hardening process: Evaluation of gear deformations and parameter optimization. (2024). Processes, 12(7), 1428. <https://doi.org/10.3390/pr12071428>
24. Optimization design of quenching and tempering parameters for crankshaft based on response surface methodology. (2024). [PMC article, PMC11312957].
25. Six Sigma model with implementation of SAPF to enhance power quality and product quality at the induction heat treatment process. (2024). [PMC article, PMC12215893].
26. Deformation mechanism and process optimization of 17Cr2Ni2Mo gear shaft heat treatment. (2025). Journal of Mechanical Science and Technology. <https://doi.org/10.1007/s12206-025-0931-9>
27. Grand View Research. (2025). India heat treating market size, share & trends report, 2024–2030.
28. Heat treatment control technology of high-strength steel gears based on support vector machine. (2025). Scientific Reports. <https://doi.org/10.1038/s41598-025-92312-1>
29. Small Industries Development Bank of India (SIDBI). (2025). Understanding the Indian MSME sector: Progress and challenges.