A CEMENTOPS AI WHITE PAPER

The Cost of Ignoring

The Cement Industry's Knowledge Exodus — and the Price of Waiting

Joshua D. Larkin | Founder & CEO, CementOps AI | jlarkin@cementops.ai

August 2026

Executive Summary

The cement industry is losing its most experienced people faster than it can replace them, and the knowledge they carry is leaving with them. This is not a forecast — retirement dates are on the calendar. The pipeline behind them is thin, and the industry's response has mostly been to hope the problem resolves itself.

This paper argues that the right way to evaluate that response is not the cost of acting — recruiting, training, knowledge capture — but the cost of ignoring (COI): the compounding operational, financial, and safety cost a plant absorbs each year it lets irreplaceable expertise walk out undocumented. Unlike most capital decisions, this one has a closing window: a program started today can interview experts who will not be available to one started in three years.

Section 3 builds an explicit expected-value model for a reference 1.2 Mtpa plant. It does not treat the full cost of a kiln outage as a knowledge cost; it estimates, channel by channel, only the portion reasonably attributable to missing experienced judgment — roughly US$196,000 to US$1.8 million a year, base case near US$625,000, compounding to US$1.1–13.1 million across a five-year delay (base case US$4.0 million). Every assumption is published so readers can substitute their own. These estimate exposure, not savings from any intervention, and the spread is an illustrative sensitivity band, not a confidence interval.

Section 4 covers the exposure no annual average can hold. Routine injury cost attributable to inexperience is modest — about US$27,000 a year here. But after an accident MSHA may issue a Section 103(k) order and control the return to normal. As a deliberate stress test, an order idling this kiln for 100 days would put roughly US$14.8 million of contribution margin at risk — several times the five-year modeled exposure. No probability is assigned to that outcome.

The response is to find the imminent single points of failure, capture one domain under expert governance, measure the operational outcome, and expand only where evidence supports it.

1. The Demographic Math Is Not in Dispute

Recent workforce studies of heavy industry point consistently in the same direction. Deloitte and The Manufacturing Institute project that U.S. manufacturing will need as many as 3.8 million new workers between 2024 and 2033 — and that 1.9 million of those jobs could go unfilled if current talent-pipeline trends hold. Deloitte and the Institute put roughly half of the projected openings at risk of going unfilled on current trends [1]. Nearly two-thirds of manufacturers already name attracting and retaining talent as their most frequently cited business challenge — ahead of supply chains, ahead of regulation, ahead of demand [2].

The age structure behind those numbers is the part that should concern cement specifically. Census Bureau workforce data illustrate how sharply manufacturing has aged: in North Carolina, one manufacturing-heavy state, workers aged 55 and older grew from roughly 12 percent of manufacturing employment in 1994 to about 30 percent in 2024, against a rise from 11 to 24 percent across all industries in the state [3]. Those datasets record age and sector, not job title. My own read as a practitioner — and it should be read as that, not as a finding of the data — is that in a cement plant the 55-plus cohort skews heavily toward kiln operators, shift supervisors, master mechanics, and QC veterans: the roles where judgment, not headcount, is the scarce input.

The cement sector's own data confirm the exposure. A 2025 survey of more than 150 senior cement industry leaders found that 80 percent of organizations are already experiencing talent gaps, and that 17 percent report being severely hindered operationally as a result [4]. Globally, Korn Ferry's talent-crunch modeling projected a manufacturing labor shortfall of 7.9 million workers by 2030 and US$607.14 billion in unrealized annual output [5].

> *80 percent of cement organizations report active talent gaps. 17 percent say they are already severely hindered operationally.*

2. Why Cement Is Uniquely Exposed

Every industrial sector faces some version of this demographic squeeze. Cement faces a harsher version of it, for reasons built into the physics and geography of the business.

The knowledge is tacit, and the ramp is long

A pyroprocessing line is not an assembly line. Competent kiln operation is pattern recognition accumulated over years: reading a coating condition from burning-zone imagery and amp trends, hearing the difference between a healthy mill and one about to trip, knowing which combination of feed, fuel, and draft moves will bring a ring down without cratering the burning zone. Little of this lives in standard operating procedures, because procedures capture what to do in defined situations — and the situations that destroy value are precisely the undefined ones.

In my own experience across twelve years in kiln operations, quality control, and reliability, the honest ramp to independent judgment in a control-room or senior maintenance role is five to ten years, and the plant leaders I speak with describe the same range. I offer that as practitioner observation rather than survey data. It is also the crux of the entire problem: when a thirty-year operator retires, the replacement is not behind by a training course; they are behind by a career.

The labor market is where the limestone is

Cement plants sit on their quarries, which means rural counties with thin labor markets. The plant cannot relocate to the talent, and the talent pool within commuting distance is small, aging, and courted by every other industrial employer in the region. Meanwhile the industry competes for young technical workers against sectors with better perceived image, urban locations, and comparable pay — and, by its own leaders' admission, has under-invested for decades in making the case for itself as a career [4].

The process is continuous, and errors are expensive

A 24/7 pyroprocess concentrates the cost of inexperience. In discrete manufacturing, a bad decision scraps parts. In cement, a bad decision during an upset can mean refractory damage, a multi-day outage, a quality excursion that reaches customers, or a reportable safety event. The scarcest resource in the industry is not clinker capacity; it is calibrated judgment available at 3 a.m. when the kiln starts telling you something is wrong.

3. The Cost of Ignoring: An Expected-Value Model

Executives are accustomed to evaluating the cost of action — a program has a price, a headcount, a line item. The cost of inaction hides in operating variance, and so it routinely escapes scrutiny.

It also invites a specific analytical error, which this section is built to avoid. The total cost of a kiln outage is not a knowledge-loss cost. Most outages would happen to some degree regardless of who is on shift, and most would cost something regardless of how they were handled. What belongs in a COI estimate is only the portion of loss that better-retained expertise would plausibly have prevented or shortened. Every channel below is therefore estimated as:

expected annual loss = exposure base × incidence × knowledge-attributable fraction

The reference plant is a single 1.2 million-tonne-per-year clinker line: 3,300 tonnes of clinker per day, 0.80 clinker factor, US$45 per tonne of cement contribution margin, 3.5 GJ per tonne specific heat consumption, US$4.00 per gigajoule fuel, US$14 million annual maintenance spend. Those are deliberately mid-range figures for a U.S. plant, not best or worst case.

3.1 What a single event actually costs

Start with the event, because it anchors the scale — then convert it to an expected annual figure, which is the number that belongs in a budget.

A 72-hour unplanned kiln outage removes roughly 9,900 tonnes of clinker. Whether that becomes lost revenue depends on the clinker silo, and this is where most published downtime arithmetic goes wrong. In a shoulder season with healthy inventory, the finish mills keep running and the plant absorbs the outage as fuel, refractory life, and fixed cost carried against zero production. In peak season on tight inventory — which is when kilns are pushed hardest and fail most often — the same outage reaches the customer as missed shipments and spot purchases.

Take the peak-season case and count contribution margin rather than revenue. At a 0.80 clinker factor, 9,900 tonnes of clinker represents roughly 12,400 tonnes of finished cement; at US$45 per tonne of contribution margin, about US$560,000. Add reheat fuel for a controlled 24-to-36-hour heat-up curve, refractory life consumed by an unplanned thermal cycle, three days of fixed cost, spot purchases or freight premiums to hold a contract, and the overtime and contractor spend of the recovery, and the full cost of a single three-day outage can approach or exceed US$1 million depending on recovery, inventory, and commercial conditions.

That is the event cost. It is not the annual knowledge-loss cost, and presenting it as such would overstate the case by roughly an order of magnitude.

3.2 From event cost to expected annual exposure

The reference plant loses about 440 hours a year to unplanned kiln downtime — roughly 5 percent of calendar hours, a mid-range figure. The question that matters is what share of those hours is caused or extended by gaps in experienced judgment rather than by equipment condition, raw materials, or bad luck. There is no published figure for that fraction in cement, and any paper claiming one should be treated with suspicion. The base case here assumes 15 percent, with a range of 10 to 22 percent, and the sensitivity of the result to that assumption is visible in the table below.

Applying the same discipline across six channels produces Table 1. Experienced judgment reduces escalation and recovery time; it does not eliminate the underlying events, and the model does not assume it does.

Loss channel	Base-case basis	Low	Base	High
Unplanned kiln downtime	440 h/yr unplanned; 15% attributable; 40% realized as lost sales	71	**283**	899
Thermal drift (SHC)	1.0% SHC deterioration; 40% attributable	21	**67**	185
Quality variability	0.30 pt clinker-factor giveaway; 45% attributable	17	**51**	135
Maintenance rework	1.8% of maintenance spend as rework; 45% attributable	42	**113**	252
Expert-concentration premium	US$130k overtime/contractor premium; 65% attributable	35	**85**	192
Routine injury exposure	1 lost-time + 2 medical-treatment-only recordables; direct + OSHA indirect; 30% attributable	10	**27**	102
Expected annual loss	US$ thousands per year, single 1.2 Mtpa	**196**	**625**	**1,765**

Loss channel	Base-case basis	Low	Base	High
	plant			

Table 1. *Expected annual knowledge-attributable loss, US$ thousands, for the reference 1.2 Mtpa plant. Base-case basis column states the incidence and attributable fraction used; low and high cases vary both. Full parameter set in the Model Assumptions appendix. The injury channel prices routine recordables only, using employer direct cost plus the OSHA Safety Pays indirect multipliers [8]. Severe low-probability outcomes are deliberately excluded from this table and are treated separately in Section 4, because they cannot honestly be reduced to an annual expected value.*

These figures estimate exposure, not savings from any particular intervention. Recoverable value depends on capture effectiveness, expert validation, workforce adoption, sustained use, and the availability of lower-cost alternatives. A program that captures most of the relevant knowledge and achieves good adoption might address a meaningful fraction of the exposure below; no vendor, including us, can responsibly promise all of it.

The base case is about US$625,000 per year at a single plant, within a range of roughly US$196,000 to US$1.8 million. The low and high columns are illustrative sensitivity scenarios, not a confidence interval: each combines every assumption at the same end of its plausible band, so the high case is a stress test rather than an expected upper bound. The range is wide because the attributable fractions are genuinely uncertain and genuinely plant-specific. That width is not a weakness in the argument; it is the argument for measuring your own plant rather than adopting anyone else's model, which is what the companion register is for.

3.3 The compounding spiral

The second-order effect is the one that turns a staffing problem into an operational crisis. Each retirement increases the load on the remaining experts: more call-outs, more overtime, more shifts covering for juniors. Overload accelerates their exits — retirement dates get pulled forward, poaching offers get accepted. Each accelerated exit increases the load further. Plants that ignore the early phase of this spiral discover that the endgame arrives faster than the demographics alone predicted, because attrition is not independent: it cascades.

The model reflects this by escalating the annual figure as the bench thins — 12 percent per year in the base case, 5 percent low, 20 percent high. That single parameter is what turns a manageable annual number into the cumulative exposure in Figure 1, and it is the assumption most worth arguing about.

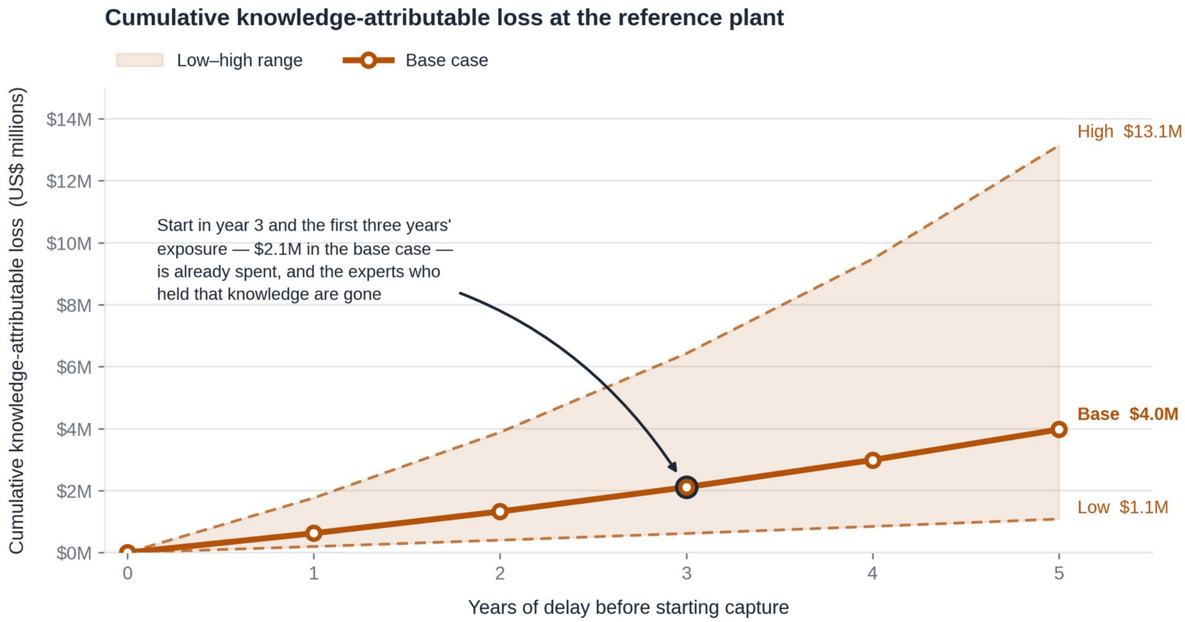

Figure 1. *Cumulative knowledge-attributable loss at the reference plant, computed from the Table 1 channels (Section 3.3 spiral) with annual escalation of 5 / 12 / 20 percent for the low, base, and high cases. Year-five cumulative: US$1.1M low, US$4.0M base, US$13.1M high. Every input is listed in the Model Assumptions appendix; the model is reproducible from that table alone.*

> *The cost of acting is a line item. The cost of ignoring is a spiral — and every year of delay removes experts that no future program can interview.*

4. Safety: The Exposure That Defeats Expected Value

The six channels in Section 3 can be modeled because they recur often enough to carry an expected value. Safety is different in kind, and it is the reason this paper does not end at Table 1.

4.1 The routine cost is modest — the hidden multiple is not

MSHA has treated tenure as a risk factor seriously enough to act on it. When the agency launched a compliance assistance initiative directed at this problem, it recorded 903 injuries to miners with one year or less experience at the mine over an eighteen-month window, and reported that seven of the eight coal mining fatalities in the year to that point involved miners with a year or less of experience at the mine [7]. Those are absolute counts rather than rates, drawn from coal rather than cement, so the defensible conclusion is not a risk multiple but the one MSHA itself drew: low mine and job tenure identifies a priority risk category. Separately and long predating that initiative, Parts 46 and 48 impose new-miner, experienced-miner, and task-training obligations intended to establish a required baseline competency.

Table 1 prices routine injury exposure at about US$27,000 a year attributable to tenure gaps, using employer direct cost grossed up by OSHA's Safety Pays estimates. OSHA puts indirect costs at 1.1

times direct cost for higher-cost injuries (US$10,000 and above) and 4.5 times direct cost for lower-cost injuries (under US$3,000), producing total direct-plus-indirect factors of 2.1 and 5.5 respectively [8]. OSHA describes these as general estimates based on limited research that can vary significantly between employers, so treat the figure as an order of magnitude. It is deliberately small: routine injuries at one plant are not where the money is.

The ratio itself, though, is worth pausing on, because the indirect burden runs inversely to severity. Investigation time, retraining, backfill, supervisory attention, and disrupted work sequences are close to fixed per event regardless of how badly someone was hurt. A strained shift absorbs a minor injury with a shrug, records the direct cost, and never sees the multiple — and the people whose hours get consumed absorbing each one are the same scarce experts already carrying the production load in Section 3.3.

4.2 The tail: when the regulator controls the return to normal

The exposure that matters most cannot be priced by frequency, because magnitude does the work instead.

Following an accident, MSHA may issue an order under Section 103(k) of the Mine Act and require approval of recovery and return-to-normal plans for affected areas. Depending on its scope, such an order can materially restrict or halt production. It does not automatically shut an entire cement plant — but the decision on when work resumes stops being the operator's.

Duration is the variable that matters. In one reported underground-coal case involving suspected heating and elevated carbon monoxide, a 103(k) order remained in effect for 100 days [9]. That is not a cement precedent and should not be read as one. As an intentionally severe stress test, however: if an order or its associated recovery process idled the reference kiln for 30 days, the modeled contribution-margin exposure would be roughly US$3.1 million once the clinker silo is drawn down; at 100 days, about US$14.8 million. The five-year base-case COI across all six channels in Table 1 is US$4.0 million.

> *A 30-day stoppage costs about three-quarters of what five full years of modeled knowledge loss costs. A 100-day stoppage costs nearly four times as much — and neither figure includes the human cost, which does not belong in a financial model.*

Two things must be said plainly about that comparison. It is a consequence calculation, not an expected value: this paper assigns no probability to a serious accident at any particular plant, and any vendor offering you one is selling something. And it prices only the regulator's response, because that is what a financial model can honestly contain.

4.3 What hazard knowledge actually walks out

What leaves with a thirty-year hand is not a private list of uncorrected hazards. Known hazards must be corrected and formally documented, and anything material carried only in someone's memory is itself a control failure. What leaves is the pattern layer sitting above the hazard register: which conditions have historically preceded a near miss, which task interfaces routinely create confusion between crews or with contractors, which abnormal signatures show up before an upset, and the operational reasons a particular control was put there in the first place.

Part 46 and Part 48 training will get a new miner to baseline competency. Neither conveys that pattern layer, because it is not a procedure — it is the residue of having been present for twenty years of things that nearly went wrong. When the last person who remembers why a control exists retires, the control survives as an unexplained inconvenience, and unexplained inconveniences are the ones that get worked around at 3 a.m.

5. Why the Traditional Playbook Falls Short

The industry is not unaware of the problem, and most operators have some version of a response: procedure libraries, shadowing programs, exit interviews, a learning-management system. These are necessary. On their own, they have consistently failed to close the gap, for structural reasons.

Procedures capture the what; the value is in the when and why. Written procedures describe defined tasks in defined conditions. Expert value concentrates in judgment under ambiguity — recognizing which situation you are actually in. Documents also go stale the day the process changes, and a binder in a control-room cabinet answers no questions at 3 a.m.

Shadowing doesn't scale on thin crews. Pairing juniors with seniors is the gold standard for transfer — and requires the one resource a short-staffed plant lacks: uncommitted expert hours. In practice, shadowing collapses into juniors watching seniors fight fires, learning fragments in whatever order emergencies dictate.

Exit interviews come too late. A career of pattern recognition cannot be downloaded in a two-week handover. By the time the retirement paperwork triggers a knowledge-transfer plan, the option to do this well expired years earlier.

The common failure mode across all of these is the same: knowledge gets captured, if at all, into formats that cannot be interrogated — static documents, tribal memory, one designated successor who then also leaves. What plants need is expertise that remains available at the point of work after the expert is gone.

6. What Good Looks Like

The technology conversation in cement has been dominated by automation and advanced process control — systems that optimize the process. The knowledge exodus demands a complementary category: systems that capture the judgment of experienced people and make it available, in context, to the people still on shift. Large language model systems grounded in a specific plant's documents, incident history, equipment quirks, and structured interviews with its veteran operators can help retrieve, structure, and deliver that captured expertise on demand — not generically, but for this kiln, this mill, this quarry.

Not all implementations are equal, and the industry's skepticism of software vendors is well earned. Operators evaluating this category should hold any solution — including ours — to a demanding standard:

- **Plant-specific grounding.** Without plant-specific grounding, generic industrial AI produces fluent answers that have never seen your kiln. The system must be built on the plant's own knowledge: its drawings, its failure history, its people's experience — captured while those people are still on payroll.
- **Built with practitioners, not just data scientists.** The capture process itself requires cement fluency. Someone who has stood at a burner floor asks different questions of a retiring kiln operator than someone who has not.
- **Delivered at the point of work.** Knowledge must be answerable from the control room, the shop floor, and the planner's desk — not archived in a repository nobody opens.
- **Auditable and honest about uncertainty.** Answers must cite their sources within the plant's own knowledge base and defer to humans on safety-critical judgment. Decision support, not decision replacement.
- **Scoped to what can actually be validated.** A vendor promising a whole plant's knowledge in one engagement is describing a roadmap. Ask which single domain will be validated first, by whom, and against what.

The governance questions to ask before any knowledge leaves the plant

Knowledge capture touches intellectual property, plant security, and the working relationship with the people whose expertise is being recorded. Those questions deserve answers before a pilot starts, not after. A serious vendor should be able to answer all seven in writing:

- **Expert validation.** Does captured knowledge require named human approval before it becomes answerable to anyone else? Nothing an AI drafts should reach a control room unreviewed.
- **Citation and version control.** Can every answer be traced to a source document or interview, with a version history showing what changed and when?
- **Access control and tenant isolation.** Is plant knowledge segregated per customer, with role-based access, and demonstrably not pooled across the vendor's client base?
- **Ownership.** Does the plant own its captured knowledge outright, including on termination, and in an exportable format?
- **Retention and deletion.** Can specific knowledge be corrected or deleted on request, and is there a defined retention policy?

- **Decision support, not control.** Is the system architecturally prevented from taking autonomous action on plant equipment?
- **Participation and recognition.** How are contributing employees involved, credited, and assured that the purpose is to preserve their contribution rather than to engineer their replacement? A capture program the workforce distrusts will produce compliant interviews and no real knowledge.

That last point is not a courtesy. The people whose judgment you most need to record are the ones most able to withhold it, and they will read the initiative's intent correctly whatever the launch memo says.

Where CementOps AI fits

Our recommended commercial starting point is a bounded 90-day pilot with named users, agreed workflows, baseline measures, and executive review. Within that pilot, and under separately agreed scope and security review, the engagement can open with a two-week discovery and capture sprint targeting three things: the prioritized knowledge register for the plant, the governance design — approval workflow, citation, versioning, and access control — and an initial domain prototype. Validation, deployment, and operational acceptance continue through the remainder of the pilot rather than concluding inside the sprint.

The company was founded by a cement practitioner with more than a decade in kiln operations, quality control, and reliability, on the view that the hard part of this problem is not the model but knowing what to ask a thirty-year kiln operator before they leave. Our copilots are advisory: they support decisions, they do not take them, and no drafted knowledge becomes active plant knowledge without named human approval. We would rather scope an engagement to what can be validated than describe a capability that is still on the roadmap.

Whatever solution an operator chooses, the criteria above are the bar.

7. A Twelve-Month Playbook

None of this requires waiting for a corporate initiative. A plant leadership team can materially reduce its knowledge-loss exposure inside a year.

1. **Build a knowledge risk register (month 1–2).** List every role and named individual whose departure would degrade operations. Score each on retirement horizon, criticality, and documentation coverage. The value of the exercise is that single points of failure become visible in one place, and in priority order, for the first time. A companion template for this step accompanies this paper.
2. **Prioritize ruthlessly (month 2).** Rank by imminence multiplied by criticality. The kiln-operations veteran retiring next spring outranks everything else on the list.

3. **Start structured capture now (months 2–6).** Standing interviews, recorded walkdowns, incident retrospectives while memories are fresh, and AI-assisted knowledge-loading for the highest-priority experts. The unit of progress is captured judgment — failure signatures, decision rules, 'here's what the gauges don't tell you' — not page count.
4. **Put the knowledge where the work is (months 4–9).** Deploy captured expertise into tools people actually consult in the flow of work: control room, maintenance planning, field execution. Adoption, not archiving, is the goal.
5. **Measure what matters (months 6–12).** Time-to-competency for new operators, first-pass startup success, repeat-incident rate, mean time to repair on knowledge-intensive failures, and overtime concentration on your named experts. These are the numbers COI shows up in — and they are how you replace this paper's assumed attributable fractions with your own measured ones.
6. **Institutionalize the trigger (month 12).** Policy: no critical expert reaches 18 months from expected departure without an active capture plan. Retirement should never again be the event that starts the conversation.

8. Conclusion: The Window Is the Point

Most industrial risks are stationary: they wait while you deliberate. This one is not. Every quarter of deliberation, more of the industry's accumulated judgment reaches a retirement date, and the option to capture it expires — permanently, because the asset in question is a person, and people do not stay available for the convenience of next year's budget.

There is no industry-wide deadline, and any paper that offers one is guessing. The actionable window is plant-specific, and it closes when your next critical expert walks out. Every operator should calculate that date rather than assume the industry shares a common schedule. For most plants that date is already on a calendar somewhere in the HR system; it has simply never been read as a financial event.

The demographic math is settled. The cement-specific exposure is documented by the industry's own leaders. The costs — outages extended, drift accumulated, injuries risked, spirals entered — are estimable and, for most plants, already being paid in forms that never appear on a line item labeled 'knowledge loss.' The question is not whether to spend on this problem. It is whether to pay deliberately, now, while the experts are still in the building — or to keep paying accidentally, at compounding rates, after they are gone.

Companion Tool: The Knowledge Risk Register

This paper is accompanied by the CementOps AI Knowledge Risk Register — a working template implementing Step 1 of the playbook. It scores each critical role on retirement horizon, operational criticality, and documentation coverage, computes a prioritized exposure ranking, and surfaces the one number worth taking to a budget conversation: how many Tier 1 roles have a departure date inside 24 months.

Request the Knowledge Risk Register · www.cementops.ai

Model Assumptions

Reference plant: 1.2 Mtpa clinker; 3,300 t/d; 0.80 clinker factor; US$45/t cement contribution margin; 3.5 GJ/t specific heat consumption; US$4.00/GJ fuel; US$14M annual maintenance spend; 8,760 calendar hours.

Low / base / high parameters — unplanned kiln downtime 260 / 440 / 700 h per year, knowledge-attributable 10 / 15 / 22 percent, margin realization 25 / 40 / 55 percent, non-margin cost US$800 / 1,200 / 1,600 per hour. Thermal drift 0.5 / 1.0 / 2.0 percent SHC, attributable 25 / 40 / 55 percent. Quality giveaway 0.15 / 0.30 / 0.60 clinker-factor points at US$375,000 per point, attributable 30 / 45 / 60 percent. Maintenance rework 1.0 / 1.8 / 3.0 percent of spend, attributable 30 / 45 / 60 percent. Expert-concentration premium US$70k / 130k / 240k, attributable 50 / 65 / 80 percent. Routine injury exposure 1 / 1 / 2 lost-time plus 1 / 2 / 3 medical-treatment-only recordables per year at direct costs of US$20k / 30k / 45k and US$1.8k / 2.5k / 4.0k respectively, grossed up by the OSHA Safety Pays indirect multipliers (2.1x and 5.5x of direct), attributable 20 / 30 / 40 percent. Annual escalation 5 / 12 / 20 percent.

Section 4 shutdown figures assume 3,288 t/d of clinker foregone with 55 percent of the resulting cement converting to lost sales over 30 days and 80 percent over 100 days, reflecting clinker-silo drawdown. Those are consequence calculations and carry no probability. These are the author's estimates, informed by practitioner experience and the cited sources. They are not measurements, and no attributable fraction in this model is drawn from published cement-industry research, because none exists. Substitute your own; the register is designed to help you derive them.

About the Author

Joshua D. Larkin is the founder and CEO of CementOps AI and a cement manufacturing practitioner with more than twelve years of hands-on experience spanning kiln operations, quality control, reliability, and plant control systems. His work on MSHA compliance and safety competency in the cement industry has been published in Cement Optimized (2026), with a second paper forthcoming. He can be reached at jlarkin@cementops.ai.

About CementOps AI

CementOps AI builds advisory AI copilots for cement operations. Engagements begin with a bounded 90-day pilot, which can open with a two-week discovery and capture sprint establishing a plant's prioritized knowledge register and governance design. Customer-specific knowledge loading, system connections, and custom workflows require separate written scope, security review, and commercial approval. Copilots support decisions rather than take them; the platform does not operate equipment, and captured knowledge becomes active only on named human approval.

References

1. Deloitte and The Manufacturing Institute, "Taking Charge: Manufacturers Support Growth with Active Workforce Strategies," April 2024. Based on surveys of 200+ U.S. manufacturers plus labor-market and economic analysis. https://www.deloitte.com/us/en/Industries/industrial-construction/articles/future-of-manufacturing-skills-gap-study.html

2. National Association of Manufacturers, Manufacturers' Outlook Survey, Q1 2024. https://nam.org/2024-first-quarter-manufacturers-outlook-survey/

3. U.S. Census Bureau, Quarterly Workforce Indicators, as analyzed by the North Carolina Department of Commerce, 2025. The QWI data record worker age by sector and do not identify occupation within manufacturing. https://www.commerce.nc.gov/news/the-lead-feed/north-carolinas-manufacturing-workforce-faces-growing-retirement-risk

4. Beaumont Bailey, "The Cement Industry's Talent Crisis," 2025. Survey of 150+ senior cement industry leaders. https://www.beaumontbailey.com/the-cement-industrys-talent-crisis/

5. Korn Ferry, "Future of Work: The Global Talent Crunch," 2018. Manufacturing sector: 7.9 million-worker global shortfall by 2030; US$607.14 billion unrealized annual output. https://www.kornferry.com/content/dam/kornferry/docs/pdfs/KF-Future-of-Work-Talent-Crunch-Report.pdf

6. Siemens, "The True Cost of Downtime 2024." Findings derive from 181 interviews with large industrial organizations, extrapolated to the world's 500 largest companies; the heavy-industry cost escalation is cited here as a directional trend and is not applied numerically to the reference plant. https://www.siemens.com/global/en/products/automation/topic-areas/value-of-reliability.html

7. U.S. Department of Labor, Mine Safety and Health Administration, compliance assistance initiative on less experienced miners, June 19, 2017. Absolute injury and fatality counts; exposure denominators not published. https://www.dol.gov/newsroom/releases/msha/msha20170619

8. Occupational Safety and Health Administration, Safety Pays: Background of Cost Estimates. Indirect-cost multipliers of 4.5x (direct cost under US$3,000) declining to 1.1x (US$10,000 and above), derived from Business Roundtable, "Improving Construction Safety Performance," and Stanford University Department of Civil Engineering research. https://www.osha.gov/safetypays/background

9. Federal Mine Safety and Health Act of 1977, Section 103(k), 30 U.S.C. 813(k); see also MSHA Accident Investigation Procedures Handbook PH20-I-4. The 100-day duration comes from a reported underground-coal case involving suspected heating and elevated carbon monoxide, described in "Section 103(k) Shutdowns," Equipment World — not a cement precedent. https://www.equipmentworld.com/workforce/safety/article/14949969/section-103k-shutdowns

---

The Knowledge Risk Register

Step 1 of the playbook in Section 7, as a working template — every critical role scored by retirement horizon, operational criticality, and documentation coverage, with a prioritized exposure ranking for your plant.

Request a copy · Joshua D. Larkin, CementOps AI · www.cementops.ai