ParagraphS
Paragraphs Environmental Design

Original Research

DOI: 10.5281/zenodo.24162475

OPEN ACCESS

***CORRESPONDENCE**
Adrian Luchini
adrianluchini87@gmail.com

RECEIVED: May 2025
ACCEPTED: May 2025
PUBLISHED: July 2025

ISSN-L 2960-2874

VOLUME 3, ISSUE 3
2025

CITATION
Luchini A. (2025). Embodied Carbon of Concrete Mixes at Comparable Strength: An Analysis of Public Environmental Product Declarations. Paragraphs Environmental Design, 3(3), 98-115. DOI: 10.5281/zenodo.24162475.

COPYRIGHT
© 2025 Luchini A. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Embodied Carbon of Concrete Mixes at Comparable Strength: An Analysis of Public Environmental Product Declarations

Adrian Luchini*

Raymond E. Maritz Professor of Architecture Emeritus, Washington University in St. Louis

Abstract

Comparing concrete environmental product declarations requires a defensible relationship between the quantity of material assessed and the performance expected from it. Compressive strength provides an accessible comparison variable, but identical strength labels do not establish identical construction requirements or service lives. This article presents an original descriptive reanalysis of two public declarations published before 1 December 2024: 48 normal-weight mixture scenarios in the National Ready Mixed Concrete Association's 2022 industry-average declaration and all 29 mixtures in the DuPont plant table of CalPortland's 2019 Washington declaration. The documents are analyzed separately because their product category rule versions and data contexts differ. The primary measure is cradle-to-gate global warming potential in kilograms of carbon dioxide equivalent per cubic metre. Within the six industry-model strength classes, the difference between the modeled zero-supplementary-cementitious-material reference and the lowest reported scenario is 36.07-39.18%, equivalent to 100.84-224.95 kg CO_2e/m^3. Conversely, differences between the 40% fly-ash and 40% slag scenarios remain below 0.1%. The producer's 4,000-psi and 5,000-psi groups exhibit substantial declared dispersion, but missing mixture-level performance information prevents interpreting those ranges as available substitution savings. Original calculations distinguish absolute carbon headroom, baseline-dependent percentage reductions, and deterministic ranking-reversal thresholds from statistical uncertainty. Dividing global warming potential by nominal strength adds no ranking information within an exact strength group and can conceal rising emissions per cubic metre across groups. The synthesis is a staged comparison method: establish accounting compatibility, establish application-specific performance eligibility, and then evaluate carbon differences at the relevant decision scale. Public declarations support useful screening decisions when their evidential limits remain attached to the results.

Keywords: concrete; embodied carbon; environmental product declaration; compressive strength; supplementary cementitious materials; comparative life cycle assessment

10.5281/zenodo.24162475

Introduction

Concrete procurement creates a recurring analytical problem: several mixtures may carry the same specified compressive strength while reporting different production-related climate impacts. The existence of such differences suggests an opportunity to reduce emissions without changing the design requirement. However, a strength label captures only part of what concrete must do, and the reported environmental difference may also depend on modeling conventions. A useful analysis must therefore answer two questions together: how large are the declared differences after strength is controlled, and what additional evidence would make those differences relevant to an actual substitution? This article addresses that combined problem through a bounded reanalysis of public environmental product declarations, or EPDs. It treats strength matching as an initial condition for comparison and examines how the interpretation changes when information on production methods, other performance properties, and uncertainty is incomplete.

The material basis for investigating these comparisons is well established. Cement and concrete research identifies clinker substitution, efficient binder use, and reductions in material demand as complementary routes to lower emissions (Habert et al., 2020; UN Environment et al., 2018). Yet those routes act at different scales. A lower-impact binder can reduce the reported burden of a cubic metre, whereas a structural redesign can change how many cubic metres are needed. Combining these effects without identifying the reference design can produce an impressive percentage with an unclear denominator. The present article concentrates on the first decision scale, the declared product volume, and uses the limitations of that scale to specify when further structural or life-cycle assessment is needed. This narrower focus makes the quantitative results easier to audit and the conclusions less dependent on untested assumptions about buildings or infrastructure.

Public EPDs offer an unusually practical connection between environmental assessment and products available to designers. They place characterized environmental indicators alongside product identities, declared quantities, and some technical properties. Third-party verification and product category rules provide a framework for disclosure, but the existence of that framework does not guarantee that two declarations contain the same performance information or use interchangeable background assumptions. ISO 14025 establishes the principles of Type III environmental declarations, while ISO 21930 provides construction-specific core rules (International Organization for Standardization [ISO], 2006, 2017). The analytical implication is that the label “EPD” identifies a category of evidence rather than a complete comparison protocol. Before using its numbers to rank alternatives, an analyst must establish what the declaration represents and which dimensions of the intended comparison remain unresolved.

There is already substantial scholarship on concrete EPDs, so originality cannot credibly rest on claiming to discover their variability. Anderson and Moncaster (2020) examined published declarations for concrete and its constituents, and Broyles et al. (2024) made a much larger compiled dataset available for subsequent analysis. Those contributions make broad comparisons possible, but increasing the number of records does not by itself resolve the difference between a modeled scenario, a catalogued product, and a verified substitute in a particular project. The contribution here is an explicit treatment of those distinctions using a small, fully traceable corpus. The article combines within-class arithmetic, a separate producer-level diagnostic, and algebraic sensitivity tests to show which conclusions follow directly from the reported values and which would require additional evidence. It offers analytical transparency rather than a claim to market coverage.

The main proposition is that matching concrete by strength can expose environmentally meaningful choices. The counterargument is that strength matching may preserve other mismatches, including age at testing, density, air entrainment, workability, and exposure resistance. These positions need not lead either to unqualified ranking or to abandoning comparison. They can be reconciled by specifying the level of equivalence actually established. A comparison may be valid as a description of declared values within a common strength class while remaining

insufficient as a purchasing recommendation. The article develops that distinction throughout its methods and interpretation. Its strongest numerical findings concern internally consistent modeled contrasts; its producer case deliberately tests what happens when the public record supports matching by strength and age but leaves other relevant properties undisclosed.

Three research questions organize the analysis. First, what absolute and relative differences in A1-A3 global warming potential occur among modeled binder scenarios within the same declared strength class? Second, how much dispersion remains among one producer's mixtures sharing a specified strength and reference age? Third, how do baseline choice, strength normalization, and sensitivity to changes in declared values alter the meaning of the resulting rankings? These questions are connected, but they are not interchangeable. The first concerns differences under a shared modeling framework, the second concerns the information available in a product catalog, and the third concerns the mathematical behavior of comparison metrics. Their synthesis supports a practical argument: a useful carbon comparison must preserve the identity of its reference, the conditions of its performance match, and the evidence supporting its numerical resolution.

The publication cutoff is 30 November 2024, inclusive, and applies to the literature and declaration editions used as evidence. The publication cutoff matters because a live repository can replace an earlier file, extend a declaration's validity, or feature a later edition whose numbers were unavailable before the cutoff. The historical documents support analysis of the values they declared, but their original validity periods have ended. The findings consequently concern the historical evidence base and a reusable analytical procedure, rather than an assertion about products currently offered or declarations currently valid for procurement.

Literature and conceptual framework

Comparable compressive strength provides a defensible starting point because strength is routinely specified and materially constrains concrete selection. A cubic metre of low-strength concrete cannot automatically substitute for a cubic metre required to meet a higher specified strength. Nevertheless, the number printed beside a mix does not completely describe its service: test age, strength definition, specimen basis, workability, exposure conditions, and construction sequence can change whether substitution is feasible. This distinction follows the broader requirement to integrate material performance into concrete LCA (Olsson et al., 2024). Equal nominal strength should consequently be understood as a conditioning rule that removes one obvious source of mismatch. It supports a narrower question about carbon differences among similarly classified products, while leaving final engineering equivalence to evidence that an environmental declaration alone may not provide.

The declared unit of one cubic metre is useful because it connects environmental accounting to concrete quantities ordered and placed. Its weakness arises when product amount is mistaken for delivered function. Konradsen et al. (2024) distinguish declared quantities from functional units and show that apparently consistent EPD terminology can conceal different forms of averaging. Applied to concrete, this suggests a two-stage comparison: first screen compatible declarations at a common strength and volume, then assess whether each alternative satisfies the intended structural and exposure requirements. Neither stage can replace the other. A fully specified structural comparison without dependable product inventories remains uncertain, while an impeccably harmonized product comparison may still select an unsuitable material. The comparison requires a chain of evidence connecting inventory compatibility, material qualification, and the service required of the completed element.

Strength-normalized indicators help expose inefficient binder use, particularly when strength gains are considered alongside the resources required to obtain them. Damineli et al. (2010) introduced binder intensity and carbon intensity as ways to relate material consumption and emissions to mechanical strength. These ratios are informative, but their interpretation depends on the decision. Dividing GWP by strength does not demonstrate that a project can reduce concrete volume in inverse proportion to compressive strength; geometric, stiffness, reinforcement, and

construction requirements also constrain design. Within an exactly matched strength class, division by the same denominator cannot change the ranking of mixes. Carbon per MPa is therefore best treated here as a secondary descriptive indicator. The principal procurement comparison remains the absolute cradle-to-gate burden for equal quantities of concrete satisfying the same specified strength condition.

Higher concrete strength can nevertheless improve environmental performance when it enables a material-efficient structural redesign. Habert et al. (2012) illustrated this relationship in bridge comparisons, where a higher burden per cubic metre could be counterbalanced by changes in material quantities at the structure level. The finding also clarifies why a product ranking and a structural ranking answer different questions. A low-GWP mix may be preferable for an unchanged element, while another mix may become preferable after a credible redesign reduces concrete or reinforcement requirements. Neither conclusion follows automatically from the other. The present strength-conditioned analysis deliberately evaluates the first decision; its implications for the second must be expressed as hypotheses requiring structural calculations. Maintaining this distinction prevents both the dismissal of high-strength concrete and the unjustified promotion of strength-normalized ratios as universal measures of sustainability.

Binder composition offers a plausible explanation for large carbon differences among concretes of similar nominal strength. The decarbonization literature identifies reduced clinker content and more efficient use of clinker as complementary opportunities rather than interchangeable measures (UN Environment et al., 2018). A lower clinker fraction can reduce emissions associated with the binder, yet a mixture requiring more total binder may surrender part of that advantage. Conversely, efficient proportioning can reduce binder demand without a conspicuous change in the advertised replacement percentage. The relevant quantity is thus the product of constituent amounts and their attributable impacts, considered together with all other included processes. An EPD showing lower GWP establishes a modeled product-level difference under its accounting rules; assigning that difference to a single mechanism requires composition data and a sufficiently controlled comparison.

Supplementary cementitious materials are consequently more informative when described by source and performance than when collapsed into a single replacement percentage. Juenger et al. (2019) emphasize that emerging SCMs require characterization and qualification, particularly as the availability of conventional materials becomes constrained. The combination of calcined clay and limestone provides another pathway: Scrivener et al. (2018) describe how coupled substitution can retain mechanical performance at substantially lower clinker contents under examined conditions. These findings support attention to binder chemistry, but they do not license universal equivalence among mixes containing the same nominal SCM fraction. Reaction rates, constituent characteristics, and curing conditions matter. An empirical EPD analysis should therefore preserve material categories where disclosed, avoid filling undisclosed recipes with assumed compositions, and interpret SCM-associated differences as conditional on the products and performance classes actually represented.

Allocation introduces a separate source of variation because the environmental burden assigned to an SCM is partly determined by the accounting relationship between coproducts. Chen et al. (2010) examine how alternative allocation procedures affect the burdens attributed to fly ash and blast-furnace slag. A material receiving little upstream burden can appear particularly attractive, but this does not mean that its physical processing, transport, or opportunity cost disappears. Nor does replacing one allocation convention with another reveal a uniquely correct measure of avoided emissions. For comparative EPD analysis, the immediate requirement is consistency and disclosure: products should be compared under compatible rules, and differences that cannot be reconciled should remain explicit limitations. A ranking conditional on an allocation convention can still guide a purchasing decision, provided it is not presented as a direct estimate of the economy-wide consequences of expanding demand.

The distinction between product choice and sector-wide mitigation becomes especially

important when low-burden constituents are supply-constrained. Directing a scarce SCM toward one concrete product may lower that product's reported impact while changing what other producers can obtain. This observation does not negate the benefits of using available SCMs; it identifies a different question from the attributional comparison represented by an EPD. Habert et al. (2020) place concrete decarbonization within a broader value chain that also includes material efficiency and construction decisions. The implication is that practical recommendations should combine favorable product selection with reduced demand for unnecessary material and qualification of additional suitable resources. A low reported footprint is evidence relevant to procurement, whereas a claim of net system-wide abatement requires a model of supply responses, substitution elsewhere, and total material quantities.

Verification strengthens confidence that an EPD follows a specified reporting process, yet it does not make all verified declarations directly comparable. AzariJafari et al. (2023) identify continuing limitations in using concrete EPDs for meaningful comparative assessment. Earlier probabilistic work also shows why shared data and methodological dependencies must be represented when uncertainty in alternatives is assessed (AzariJafari et al., 2021). A common background factor can move two product results together, so independently perturbing every value may distort uncertainty in their difference. Conversely, an exact numerical ranking may appear more decisive than the underlying inventories justify. The analysis should distinguish reported differences from robust decisions: inspect compatibility first, retain traceable arithmetic, and state what magnitude or direction of unresolved change would be needed to overturn the preferred option.

Durability broadens the environmental question because a material that requires earlier replacement may consume more resources over the service period. Miller (2020) demonstrates, through dynamic material-flow modeling, that longer in-use residence can reduce demand and emissions. That result supports attention to longevity, but it cannot establish the service life of a particular mix from its compressive strength or binder label. A declaration covering production only also lacks the information needed to calculate future maintenance, repair, and replacement reliably. Equal-strength screening should therefore be accompanied by explicit exposure and durability qualification, with whole-life comparisons undertaken where defensible scenarios are available. This preserves the usefulness of A1-A3 results while preventing their extension into claims that the lowest production burden necessarily produces the lowest lifetime burden.

Carbonation illustrates why a scientifically real process can still be misapplied in product comparisons. Xi et al. (2016) estimate substantial historical uptake by cement materials, but their reported offset excludes fossil-fuel emissions and concerns accumulated material stocks. Van Roijen et al. (2024) further demonstrate that delayed uptake and its timing alter the climate benefit inferred from a simple mass balance. Neither study supports deducting a universal percentage from every concrete EPD. For an analysis restricted to A1-A3, future uptake should remain outside the reported comparison; an expanded assessment would need consistent exposure, geometry, service-life, end-of-life, and temporal assumptions. The resulting position recognizes carbonation without allowing an uncertain later removal to obscure an immediate difference in production emissions, and keeps the article's conclusions aligned with the boundary actually analyzed.

Materials and methods

Research design and selection of declarations

The study uses a purposive two-document case design. The first case was selected because it presents a complete factorial portfolio of normal-weight strength classes and binder scenarios under one calculation framework. The second was selected because a complete single-plant table presents several mixtures at repeated strengths with an explicit reference age. Selecting for these features improves internal interpretability, but it does not produce a random or comprehensive sample of the concrete market. Accordingly, the analysis estimates differences within the selected declarations and does not estimate national mean emissions, market

percentiles, or the prevalence of low-carbon products. The two cases answer related questions through different evidence structures, and their results remain separate throughout. This separation follows from the differences in the evidence represented by the two sources.

The first source is the National Ready Mixed Concrete Association's industry-average EPD, issued on 3 January 2022 under identifier EPD10294. The historical PDF reports 72 scenarios, of which 48 are normal-weight products arranged as six strength groups with eight binder scenarios each; the 24 lightweight scenarios are outside the main analysis. Results use the upper endpoint of each strength range and the lower endpoint of the relevant supplementary cementitious material range. The comparison therefore concerns modeled endpoints, not measured strengths of independently sampled batches. The eight scenario labels are a modeled zero-SCM reference, 20%, 30%, and 40% fly ash, 30%, 40%, and 50% slag, and a ternary blend with 20% fly ash plus 30% slag. The reference calculation must not be misdescribed as the observed average of all products in its broader low-SCM classification (National Ready Mixed Concrete Association [NRMCA], 2022).

The second source is CalPortland Company's Washington facilities declaration, issued on 18 June 2019 under identifier NRMCAEPD: 20026. All 29 entries in its DuPont plant Table 7b are retained, rather than selecting only the first page or a convenient low-carbon pair. They report 28-day strengths in six groups, from 3,000 to 6,000 psi. The declaration provides product-specific environmental results, but the extracted table does not disclose each mixture's density class, air content, SCM percentage, or exposure requirements. Its case therefore describes dispersion among products carrying the same strength-age label. It cannot establish that all entries are normal-weight concrete or that the minimum and maximum entries are technically interchangeable. Preserving this limitation allows the producer case to test the adequacy of public comparison metadata without turning missing information into an assumed match (CalPortland Company, 2019).

The documents share an A1-A3 reporting boundary and a North American assessment context, which makes their respective within-document comparisons useful. Both apply a recovered-material convention that limits burdens assigned to SCMs to the treatment and transport required for their use; the reported comparisons remain conditional on that convention. However, the NRMCA declaration identifies the 2021 version 2.1 concrete PCR, whereas CalPortland identifies the February 2019 version 1, and their foreground data and background inventories differ. These similarities are therefore insufficient to justify a pooled ranking across the two documents. No claim is made that a producer entry performs better or worse than the industry model simply because its reported number lies above or below a model value. Instead, the study compares the form and interpretability of their internal variation. Keeping the cases separate preserves what each source independently supports.

Data extraction and comparability checks

Extraction retained the declaration identity, issue date, mixture identifier, strength information, reference age, reported GWP, declared volume unit, and exact table and PDF page. The NRMCA data come from Tables 6a-11a on physical PDF pages 15, 17, 19, 21, 23, and 25; the DuPont data come from Table 7b on physical pages 26-28. Both sources supply results per cubic metre, so the analysis uses those values directly. This avoids introducing conversion rounding through the corresponding cubic-yard tables. The NRMCA tables also contain summary minimum and maximum columns, which are excluded as duplicate summaries rather than counted as additional mixtures. The resulting datasets contain 48 and 29 unique entries. A separate extraction audit checks the GWP row against the source columns, because an apparently plausible numerical range does not establish that the correct product labels were attached to it.

Strength is retained in its source form before any secondary normalization is applied. For the industry case, the six upper-endpoint labels are 2,500, 3,000, 4,000, 5,000, 6,000, and 8,000 psi, and the underlying ranges remain visible in the data file. For the producer case, matching uses the exact stated psi value and the common 28-day age. These labels are not converted into assumed

measured mean strengths, nor are undisclosed specimen characteristics supplied from general expectations about regional practice. The secondary GWP-per-strength calculation converts psi consistently using 1 psi = 0.006894757293168 MPa. Its denominator is consequently the declared or modeled strength label, never a new experimental measurement. This preserves a clear separation between unit harmonization, which is arithmetic, and performance harmonization, which requires engineering evidence.

Accounting and technical comparability are checked separately. Accounting checks cover the declared quantity, module boundary, characterization method, PCR edition, and whether the values share a modeling context. Technical checks cover the strength-age pair, density classification where available, and disclosed requirements that might change the mixture's suitability. A product may pass the first set while remaining unresolved on the second. Treating those outcomes separately prevents a complete environmental table from creating a false impression of complete performance information. It also makes the unresolved items actionable: the next investigation can request the missing mixture specification or test evidence, rather than merely searching for another GWP value. Table 1 summarizes the evidence status of the two cases and the conclusions permitted at their respective comparison levels.

Table 1
Evidence available for the two separate comparisons

Attribute	Industry model	Producer case
Source and issue	NRMCA 2022, EPD10294	CalPortland 2019, NRMCAEPD 20026
Analytical subset	48 scenarios; six strength classes × eight binder scenarios	29 entries; complete DuPont Table 7b
Strength and age	Modeled upper class endpoints; 28 days	Declared strength labels; 28 days
Density and binder	Normal-weight; modeled SCM percentages known	Mixture-level density, air and SCM percentages unavailable
Accounting	A1-A3; kg CO_2e/m^3; TRACI 2.1; concrete PCR v2.1	A1-A3; kg CO_2e/m^3; TRACI 2.1; concrete PCR v1
Permitted inference	Within-model strength-class contrasts	Within-label catalog dispersion
Unresolved inference	Actual batch or project substitution performance	Technical equivalence and ingredient attribution

Note. Source facts: NRMCA (2022) and CalPortland Company (2019). The comparison levels are the present study's analytical interpretation. Results are not pooled.

Comparison measures and sensitivity calculations

For a matched comparison, B denotes the higher reference GWP and L denotes the lower alternative GWP, both in kg CO_2e/m^3. The absolute difference H is B minus L, and the proportional reduction R divides that difference by B. Reporting both measures is necessary because the percentage describes the change relative to a chosen reference while the absolute difference describes the product-stage carbon difference for a fixed volume. Neither quantity is a measured project saving unless the replacement is feasible and the quantity is actually procured. In the industry case, the primary reference is explicitly the modeled zero-SCM scenario in the same class. In the producer case, the corresponding calculation is described as a maximum-to-minimum catalog spread. The shared arithmetic does not erase the difference in evidential meaning between the two cases.

$$H = B - L, \quad R = \frac{B - L}{B} \times 100\%$$

The choice of reference is examined directly rather than hidden in a single headline percentage. An existing 20% fly-ash scenario may be a more relevant starting point than zero SCM for some projects, but the present sources do not establish which scenario any purchaser would otherwise use. A second calculation therefore changes only the analytical reference, comparing the ternary scenario with the 20% fly-ash scenario within each industry strength class. This is a baseline sensitivity analysis, not a claim about market practice. By keeping the final option constant and changing the starting point, it shows how a technically correct percentage can communicate different amounts of incremental improvement. The interpretation follows the counterfactual that has actually been defined, rather than treating the largest available reduction as a universal opportunity.

The strength-normalized indicator I divides GWP G by the strength label f. It is included because intensity indicators can help identify patterns in the use of binder and in material performance (Damineli et al., 2010). However, when all products in a comparison have the same f, normalization cannot change their ordering or their relative differences. Across strength classes it introduces a different denominator and may alter rankings without showing a reduction in product volume or total emissions. The article therefore presents I only as a diagnostic of metric behavior. It does not use it to select a structurally preferred mixture or to convert unlike strength classes into equivalent functions. A project-level comparison would need a reference structure and a calculation of the quantities actually needed to satisfy its design requirements.

$$I = \frac{G}{f}$$

Ranking sensitivity is expressed as a break-even condition rather than an invented probability distribution. Suppose the higher declared value decreases proportionally by u while the lower value increases by the same proportion. The values become equal when B times one minus u equals L times one plus u, yielding the threshold u*. A small threshold indicates that a small oppositely directed change could remove the numerical ranking; a larger threshold requires a larger change under that particular construction. The exercise does not estimate the likelihood of such changes, and u* is not a confidence interval or an EPD error rate. Its purpose is to distinguish substantial separations from fragile numerical orderings using only the reported values. Probabilistic claims would require additional information about distributions, dependencies, and the source of each uncertainty (AzariJafari et al., 2021).

$$B(1 - u) = L(1 + u), \quad u^* = \frac{B - L}{B + L}$$

The direction and dependence of changes are essential to interpreting this threshold. A common positive multiplicative adjustment to both values preserves their ranking, whereas an equal additive adjustment cancels from their absolute difference. Opposite proportional changes are therefore a diagnostic stress case, not a realistic universal model of shared inventory error. A separate use of H asks how much net additional burden attributable only to the lower-GWP option would exactly remove its initial advantage. That additional burden could arise from a verified change in delivery or construction requirements, but none is assumed to occur here. This formulation converts the declared difference into an explicit break-even quantity while keeping the missing project evidence visible. It also avoids treating a favorable A1-A3 result as permission to ignore emissions occurring outside the assessed boundary.

Statistical scope and reproducibility

All reported results are descriptive calculations on fixed declared values. The NRMCA scenarios are outputs of a common modeling structure, and the producer entries share a plant and assessment framework. Neither dataset supplies independent random observations of an underlying industry distribution or replicated measurements from which sampling error can be estimated. Accordingly, the study reports no significance tests, bootstrap confidence intervals, or regression coefficients interpreted as causal effects of SCM

content. The absence of those statistics does not prevent exact statements about the selected tables; it limits the population and causal claims that can be attached to them. Re-running the arithmetic reproduces the article's results, while evaluating their external applicability remains a separate research task.

Results

Carbon differences within modeled strength classes

The industry-model portfolio shows lower declared A1-A3 GWP for every included SCM scenario than for the modeled zero-SCM reference in the same strength class. The largest within-class separation ranges from 100.84 kg CO_2e/m^3 at the 2,500-psi endpoint to 224.95 kg CO_2e/m^3 at the 8,000-psi endpoint, corresponding to reductions of 36.07% and 39.18%, respectively. The ternary scenario is the minimum in each class. These results establish a substantial difference within the selected modeling framework, but they do not establish a universal reduction for concretes carrying the same nominal replacement percentage. The model's mixture-proportioning and inventory assumptions remain part of every comparison. The defensible finding is that strength class alone leaves substantial room for variation in declared product-stage climate impact within this coherent scenario portfolio.

Table 2

Industry-model A1-A3 GWP at each modeled strength endpoint

psi	0% SCM	20% FA	30% FA	40% FA	30% SL	40% SL	50% SL	20% FA + 30% SL
2,500	279.56	241.63	221.18	199.65	219.66	199.80	179.86	178.72
3,000	311.39	268.41	245.16	220.73	243.44	220.91	198.37	197.09
4,000	383.60	329.16	299.71	268.75	297.54	268.97	240.33	238.71
5,000	468.89	401.04	364.37	325.65	361.67	325.93	290.12	288.10
6,000	493.67	421.99	383.24	342.33	380.40	342.64	304.88	302.75
8,000	574.10	489.67	444.03	395.77	440.68	396.14	351.67	349.15

Note. All GWP values are kg CO_2e/m^3. FA = fly ash; SL = slag. Rows are endpoint labels for the original strength ranges, not measured strengths. Source: NRMCA (2022), Tables 6a-11a, PDF pp. 15, 17, 19, 21, 23 and 25. Columns represent modeled percentages.

At the 4,000-psi endpoint, the zero-SCM reference reports 383.60 kg CO_2e/m^3 and the ternary scenario reports 238.71 kg CO_2e/m^3. Their difference is 144.89 kg CO_2e/m^3, or 37.77% of the reference. The 50% slag scenario is close to the minimum at 240.33 kg CO_2e/m^3. Consequently, the broad contrast between the reference and the two lowest scenarios is much larger than the contrast between those scenarios themselves. A decision that treats all three pairwise differences as equally informative would ignore this difference in resolution. The interpretation should separate evidence for moving away from a high-impact reference from evidence for selecting the numerically lowest member of an already low-impact group. The first comparison has a substantial declared margin; the second requires closer scrutiny of uncertainty and performance constraints.

Percentage reductions increase modestly across the six endpoint classes, while absolute differences increase much more strongly. This distinction matters for quantity-based decisions: a percentage describes relative improvement, whereas multiplying a validated per-volume difference by an actual project quantity gives the corresponding product-stage difference. However, the study does not supply the quantities, mix availability, or approved substitutions for a real project. It therefore reports the per-volume headroom rather than inventing a project total or an industry-wide saving. Table 3 retains all six

classes so that the reader can see how the relative and absolute measures change together. The observed pattern supports class-specific assessment; it does not justify replacing the full set with a single average percentage and applying it to unrelated concrete demand.

Table 3
Original within-class contrasts against the modeled zero-SCM reference

psi	Reference GWP	Lowest GWP	Headroom H	Reduction %	Threshold u* %
2,500	279.56	178.72	100.84	36.07	22.00
3,000	311.39	197.09	114.30	36.71	22.48
4,000	383.60	238.71	144.89	37.77	23.28
5,000	468.89	288.10	180.79	38.56	23.88
6,000	493.67	302.75	190.92	38.67	23.97
8,000	574.10	349.15	224.95	39.18	24.37

Note. GWP and H are kg CO_2e/m^3. The lowest scenario is 20% FA + 30% SL in every row. u* is the deterministic opposite-perturbation threshold, not a confidence limit. Calculated from Table 2.

Reference choice and small differences between SCM scenarios

Changing the reference changes the stated opportunity even when the proposed lower-GWP scenario stays the same. At 4,000 psi, comparison with the 20% fly-ash scenario of 329.16 kg CO_2e/m^3 gives a ternary difference of 90.45 kg CO_2e/m^3, or 27.48%, instead of the 37.77% obtained from the zero-SCM reference. Both calculations are correct, but they answer different questions about what would otherwise be purchased or modeled. The source does not identify that counterfactual for a particular project. A defensible claim of improvement must therefore name the starting mixture and its strength class alongside the percentage. This is more than a reporting preference: choosing an unrepresentative high-impact reference can exaggerate the incremental benefit available to a purchaser whose existing specification already uses SCMs.

The modeled difference between 40% fly ash and 40% slag is only 0.15-0.37 kg CO_2e/m^3 across the strength classes, less than 0.1% relative to the slag value. The difference between the ternary scenario and 50% slag is 1.14-2.52 kg CO_2e/m^3, remaining below 0.72%. These are exact differences between the printed values, but printing to two decimal places does not demonstrate accuracy at that resolution. The data therefore support describing the pairs as numerically close under the adopted framework, rather than promoting one binder option as categorically superior. The distinction is especially relevant because the same portfolio supports much larger reductions against its high-impact reference. Preserving the near-ties strengthens the analysis by locating where the evidence ceases to support fine discrimination.

Binder labels alone also fail to order all options consistently across replacement levels. At 4,000 psi, the 30% slag scenario has slightly lower declared GWP than 30% fly ash, at 297.54 versus 299.71 kg CO_2e/m^3. Yet 40% fly ash, at 268.75 kg CO_2e/m^3, is lower than either 30% scenario. This example does not isolate material chemistry, because the complete mixture proportions and model assumptions determine each result. It does show why a general rule favoring one SCM by name cannot replace comparison of specified mixtures. The analytical unit must remain the mixture scenario at its stated strength and replacement level. A material category can guide investigation, but the observed GWP belongs to the whole modeled product rather than to a label detached from dosage and performance conditions.

Strength normalization and the denominator problem

The zero-SCM series demonstrates how a strength-normalized index can improve while emissions per cubic metre increase. Between the 2,500-psi and 8,000-psi endpoints, declared GWP rises from 279.56 to 574.10 kg CO_2e/m^3. Over the same comparison, GWP divided by the converted strength label falls from approximately 16.22 to 10.41 kg CO_2e per cubic metre per MPa. Neither result contradicts the other: strength has increased proportionally more than GWP. However, the lower ratio does not establish lower carbon for a structure, because the data do not establish a compensating reduction in concrete volume. The index is therefore informative about the relationship between two declared quantities, while the structural implication remains conditional on a separate design calculation.

The denominator problem is general rather than peculiar to these six values. If a hypothetical relationship consists of a positive fixed contribution plus a term proportional to strength, dividing by strength creates a declining fixed-contribution term even if no technology or material efficiency changes. That algebra is an explanatory example, not a fitted model of the EPD. It demonstrates why a downward curve of GWP per MPa cannot, by itself, prove improved engineering efficiency. Within a fixed strength group the ratio is redundant; across groups its relevance depends on how strength changes the required structure. The article consequently retains kg CO_2e/m^3 as its primary measure and interprets strength as a condition of comparison, rather than assuming that multiplication by strength produces a complete measure of useful construction service.

The portfolio also contains a cross-class reversal in absolute GWP: the 8,000-psi ternary scenario reports 349.15 kg CO_2e/m^3, below the 4,000-psi zero-SCM scenario's 383.60 kg CO_2e/m^3. This observation disproves a simple ordering in which every higher-strength scenario must have higher declared carbon per volume. It does not identify an approved substitution, and cross-class changes in air-entraining assumptions further limit a purely strength-based interpretation. Its useful implication is narrower: strength is an incomplete predictor of embodied carbon, even within a single portfolio. Designers should investigate the composition and assessment context of suitable options rather than assigning one universal emission factor to each strength value. Actual selection still requires the candidate to satisfy the complete specification for the intended element.

Dispersion in the producer declaration

The DuPont case contains repeated 28-day strength labels but very different amounts of GWP dispersion. The three 3,000-psi entries range from 326.85 to 332.05 kg CO_2e/m^3, a maximum-relative spread of 1.57%. The ten 4,000-psi entries range from 266.71 to 439.94 kg CO_2e/m^3, a spread of 39.38%, while the ten 5,000-psi entries range from 355.26 to 520.37 kg CO_2e/m^3, a spread of 31.73%. These are catalog ranges within a common plant declaration, not confidence limits or experimentally demonstrated emissions reductions. Their different widths show that a nominal strength label does not determine a unique reported GWP. The missing mixture-level metadata prevents identifying how much of that dispersion would remain after a complete application-specific performance match.

Table 4
Declared GWP dispersion within DuPont 28-day strength groups

psi	Entries	Minimum	Median	Maximum	Spread %
3,000	3	326.85	329.570	332.05	1.57
3,500	2	332.87	384.575	436.28	23.70
4,000	10	266.71	389.425	439.94	39.38
4,500	2	417.55	445.305	473.06	11.73
5,000	10	355.26	436.875	520.37	31.73
6,000	2	468.46	493.610	518.76	9.70

Note. GWP values are kg CO_2e/m^3. Spread = 100 × (maximum - minimum)/maximum. Medians are unweighted catalog summaries; two middle values are averaged where needed. Source: CalPortland Company (2019), Table 7b, PDF pp. 26-28. These ranges do not establish substitution savings.

The lowest and highest 4,000-psi values belong to mixture identifiers 64335 and 3324; the corresponding 5,000-psi identifiers are 63066 and 61517. Reporting those identities makes the comparison traceable and allows a subsequent investigator to request the missing specifications. However, a ranking of their published GWP values is not enough to conclude that one can replace the other. The EPD's general material list allows ingredients whose use may change density or other properties, and no mixture-level classification was inferred from the identifiers. The practical result is a focused information request: determine whether the lower-GWP entries satisfy the same application requirements before interpreting the observed spread as an achievable saving. This preserves the screening value of the data while preventing an unsupported recommendation about any particular CalPortland product.

The producer case supports the industry-model finding only at a broad descriptive level. Both show that strength labels coexist with environmental variation. They do not independently reproduce the same SCM effect: the industry case specifies modeled binder scenarios, whereas the producer table does not provide the ingredient detail required for attribution. Combining them into a larger numerical sample would obscure that distinction and create an artificial impression of replication. Keeping them separate instead identifies two different research needs. A coherent model can reveal the consequences of its own assumptions and the scale of its scenario differences; a public product catalog can reveal where additional performance disclosure is needed. Their joint contribution concerns the design of credible comparison, rather than a pooled estimate of concrete decarbonization potential.

Ranking sensitivity and carbon headroom

For the 4,000-psi industry comparison between 383.60 and 238.71 kg CO_2e/m^3, the symmetric ranking-reversal threshold is 23.28%. Under the defined stress case, one value must fall and the other rise by that proportion before their ordering disappears. For 40% slag versus 40% fly ash in the same class, the threshold is only about 0.0409%. The difference between these thresholds formalizes the difference between a broad separation and a near-tie, but it does not establish the probability that either ranking is correct. Real uncertainties may be shared, differently sized, or concentrated in particular inputs. The thresholds should therefore guide the intensity of further investigation: tiny margins are poor grounds for strong preference claims, while large margins justify examining whether technical and project conditions preserve the apparent advantage.

Table 5
Sensitivity and baseline contrasts within the 4,000-psi model class

Comparison	B	L	H	Reduction %	u* %
Zero SCM vs. ternary	383.60	238.71	144.89	37.7711	23.2826
20% FA vs. ternary	329.16	238.71	90.45	27.4790	15.9279
50% SL vs. ternary	240.33	238.71	1.62	0.6741	0.3382
40% SL vs. 40% FA	268.97	268.75	0.22	0.0818	0.0409

Note. B, L and H are kg CO_2e/m^3. Ternary = 20% FA + 30% SL. u* = (B - L)/(B + L), expressed as a percentage. Additional decimal places display arithmetic for near-ties and do not imply measured accuracy. Calculated from NRMCA (2022).

Absolute carbon headroom gives the same comparison a different practical interpretation. At 4,000 psi, the ternary scenario has 144.89 kg CO_2e/m^3 of product-stage headroom relative to zero SCM, but only 1.62 kg CO_2e/m^3 relative to 50% slag. If an otherwise eligible alternative creates a verified net additional burden outside A1-A3, that burden can be compared with the relevant headroom on a consistent quantity basis. No such burden is estimated in this study, and the headroom is not a license to omit construction emissions or assume equivalent durability. Its value is as a break-even question grounded in an observed declaration difference. A project can then investigate whether additional transport, curing, or other requirements are likely to consume that margin, using its own conditions rather than a universal adjustment factor.

Discussion

Agreement in strength and reference age provides a useful first comparison, although it cannot establish complete equivalence. Earlier concrete EPD research proposed additional performance metrics precisely because structural and pavement applications require information beyond compressive strength (Azarijafari, 2020; AzariJafari et al., 2023). This study's incomplete producer metadata makes that distinction operational: missing attributes narrow what may be concluded, even when the GWP difference is large. The synthesis is a staged decision process in which strength and age identify a candidate set, performance requirements establish eligibility, and comparable environmental information distinguishes the eligible alternatives. A low carbon score enters the final decision only after the mix's intended function has been demonstrated.

Procurement can apply these findings by making carbon performance conditional on the complete technical specification. The EPA's December 2022 interim determination already connected comparison to similar products meeting the same functional requirements and favored facility- and supply-chain-specific EPD information where available (U.S. Environmental Protection Agency, 2022). However, its historical percentile framework cannot be reproduced from the unweighted mixture portfolios studied here. A manufacturer that publishes more variants would otherwise acquire disproportionate influence on the apparent market distribution. These data are therefore suitable for identifying and evaluating candidate mixtures, while jurisdictional thresholds require an independently defined, transparent and geographically relevant reference population.

Project accounting should also weight eligible mixture impacts by the quantities actually required. If each approved mixture has volume Vj and compatible declared intensity Gj, the product-stage total is the sum of VjGj; a comparison uses the same project scope and justified quantity assumptions. The arithmetic is straightforward, but substituting a lower score while silently changing structural quantities would conceal a second design decision. Likewise, delivery and placement remain additional modules rather than disappearing because an EPD reports A1-A3. The resulting procurement record should connect each carbon calculation to a qualified mixture, a physical quantity and a declared

boundary, allowing reviewers to distinguish a material change from a redesign.

Performance matching as a documented decision

The staged comparison proposed here makes the decision record as important as the final carbon number. A project first defines the service required of the concrete, including the strength-age pair and any other conditions that affect eligibility. The environmental comparison then operates within that qualified set. This sequence might appear to restrict the opportunity to reduce carbon, especially when the public minimum is excluded, but it prevents a misleading comparison against a product that cannot perform the required task. It also makes improvement measurable: a supplier can seek qualification of a lower-impact candidate for a defined use instead of claiming superiority across all applications. The resulting record should show the original specification, the candidate's evidence of compliance, the compatible environmental declaration, and the quantity basis used to calculate the difference.

Performance-based selection should nevertheless allow genuine redesign of unnecessary constraints. A specification may contain a prescriptive ingredient requirement whose original purpose could be met by another verified route. Removing that requirement without examining its function would be unjustified, but retaining it automatically could exclude a lower-impact solution that satisfies the actual performance need. The EPD analysis cannot adjudicate such cases because it contains neither the project specification nor the supporting material tests. It can identify where investigation might be valuable. A technically reviewed change to the eligible set should be documented as a specification or design change, after which the carbon comparison can be repeated. This preserves the possibility of innovation while keeping the environmental claim attached to the conditions under which the alternative became acceptable.

Reference age deserves particular attention because it connects binder behavior to construction scheduling. SCM performance depends on material characteristics and reaction development, so the suitability of a mixture cannot be inferred from replacement percentage alone (Juenger et al., 2019). A later acceptance age might permit a different formulation, but the present 28-day comparisons do not quantify that opportunity. They also do not show whether stripping, loading, finishing, or curing requirements would change. Consequently, extending the acceptance age is a new design condition requiring its own evidence, rather than an adjustment that can be silently applied to the existing results. A future comparison should hold the project's actual time-dependent requirements explicit and measure the implications of meeting them. That would connect environmental performance with construction feasibility more directly than a single compressive-strength label.

The results also indicate how to handle close alternatives when neither clearly dominates the full decision. A difference below 1% in the modeled GWP does not establish environmental equivalence, but it provides a weak basis for a strong purchasing preference when important performance or inventory uncertainties remain. The appropriate response is to preserve the reported numbers and identify the ranking as unresolved at the available evidential resolution. Additional data could still distinguish the alternatives, particularly if the small difference is stable under a validated common model. Until then, the selection can be made on demonstrated eligibility and other relevant project criteria without falsely describing a tiny calculated lead as a reliable carbon advantage. This approach avoids both arbitrary rounding that erases the evidence and exaggerated precision that overstates it.

What total EPD values can and cannot explain

The GWP of a concrete mixture aggregates contributions from constituent production, transport, and manufacturing processes. A lower total is consistent with several possible mechanisms, including lower cement demand, lower-impact cement supply, different SCM quantities, and changes in other inputs. However, a total alone cannot uniquely identify those contributions. Many different combinations of constituent quantities and emission factors can produce the same aggregate result. Regressing the 48 scenario totals against a few replacement labels would therefore risk attributing the consequences

of joint model assumptions to a single apparent driver. The justified interpretation remains at the scenario level unless component quantities and compatible inventory factors are available. A more detailed decomposition would be valuable, but its added explanatory power must come from additional evidence rather than from mathematical fitting to already aggregated values.

Model transparency helps distinguish an internally coherent conclusion from a general empirical law. The NRMCA assessment assumes fly-ash reactivity of 75% relative to portland cement and slag reactivity of 100%; its air-entraining assumption changes above 5,000 psi. Those choices explain why cross-class trends cannot be assigned solely to strength, although within-class comparisons retain a common setting (NRMCA, 2022). Their disclosure is useful precisely because it reveals conditions that could alter the answer. The appropriate next step is not to treat the chosen reactivities as universal material constants, but to examine whether locally qualified mixtures and inventories produce similar differences. The present analysis consequently supports a conditional statement about the reported model, and it specifies the information needed to assess how far that statement transfers to another production context.

An audit trail also changes how inconsistencies should be handled. The historical NRMCA classification table contains inconsistent printed slag-range labels in its first strength group, while the product names and modeled percentages identify the scenarios used in the results tables. Ignoring that discrepancy could create a misleading reconstructed classification; inventing a corrected range would add information the source does not establish. The extraction therefore preserves the modeled scenario identity and flags the range-label issue in the audit. More generally, an apparent typographic error should be distinguished from a numerical extraction error and from a methodological ambiguity. Each requires a different response: source clarification, corrected transcription, or a restricted interpretation. An environmental dataset becomes more reliable when those distinctions remain visible instead of being absorbed into undocumented cleaning decisions.

Implications beyond the selected binder portfolio

The portfolio's focus on conventional fly ash and slag should not define the complete technological opportunity. Research on calcined clay-limestone cement identifies a technically important route to lower clinker content, while work on alkali-activated materials describes additional possibilities and application-specific constraints (Scrivener et al., 2018; Provis, 2018). Neither technology is represented in the numerical sample. Its absence cannot be used to infer inferior carbon performance, limited feasibility, or lack of relevance. Conversely, literature-level promise cannot be inserted into these tables as if it were an observed EPD result. A defensible extension would add separately qualified declarations with matched performance information and compatible accounting, then state which comparisons can be harmonized. This keeps the study open to technological change without overstating what a historical two-document corpus contains.

The same principle applies to environmental outcomes beyond climate change. GWP is the chosen endpoint because the research question concerns embodied carbon, but a lower GWP score does not describe water demand, air-pollution burdens, or every other impact category. Miller and Myers (2020) show that alternative binder assessments can involve different patterns across environmental indicators. It would therefore be inappropriate to rename the minimum-GWP scenario the universally most sustainable mixture. At the same time, the absence of a multi-impact ranking does not invalidate the climate comparison that has actually been made. The assessment should state the endpoint precisely and conduct additional category-specific checks where the project's environmental objectives require them. Any aggregation of different impacts would need its own weighting rationale, rather than being concealed within a carbon-based ranking.

Limitations and priorities for further research

The most immediate limitation is the purposive selection of two documents. It provides complete and auditable subsets suited to the research questions, but public availability, surviving historical links, and readable performance tables influence which products enter the analysis. Those

influences could exclude producers or mixtures with different environmental characteristics, and their direction cannot be estimated from the selected sample. The conclusions therefore concern the demonstrated possibilities and interpretive limits of these cases. A larger follow-up should define its population in advance, record unavailable or excluded documents, distinguish superseded editions from unique products, and disclose the consequences of missing metadata. The existing large dataset of Broyles et al. (2024) offers useful infrastructure for that work, but a wider sample would still need an explicit account of what its records represent.

The distinction between a product count and a production distribution creates another limitation. A producer may disclose many closely related mixtures while another discloses only a few broad categories. Counting every row equally then gives the first producer more influence over an apparent average or percentile, even if its delivered volume is smaller. The present study avoids market inference and reports only within-case summaries, but that restriction also means it cannot supply a representative benchmark for procurement policy. A future benchmark must define whether it describes available products, producing facilities, or concrete volumes delivered during a specified period. Each is a legitimate object of analysis, yet their statistical weights and interpretations differ. Explicitly defining that object is a prerequisite for any claim about typical performance or the best-performing portion of a market.

Temporal representativeness is also constrained. A declaration's issue date identifies a publication event, while its production data can represent an earlier period; the CalPortland assessment identifies 2017 foreground data despite its 2019 issue date. Historical results are therefore unsuitable as unqualified estimates of later plant performance. Nevertheless, they remain valid evidence for analyzing what the published declaration reported and how its comparison metrics behave. The article preserves that distinction by fixing the editions. A future trend analysis would need matched observations across time, consistent accounting, and an explanation of changes in energy, cement supply, and product composition. Comparing unrelated editions or current and historical repository averages would not, by itself, establish a real improvement in production.

Whole-life performance remains outside the empirical boundary. The article neither estimates future maintenance and replacement nor deducts later carbonation from A1-A3 emissions. This prevents unsupported whole-life claims, but it also means a project requiring those claims must undertake additional assessment. Service-life research and time-sensitive carbonation studies show why those extensions require explicit scenarios rather than a universal adjustment (Miller, 2020; Van Roijen et al., 2024). Future work should compare qualified alternatives within the same structural function, exposure conditions, reference period, and end-of-life assumptions. That would allow the product-stage advantage to be evaluated alongside downstream differences. The current results can supply one input to such an analysis; they cannot determine its answer in advance.

A productive research extension would combine a public EPD comparison with a documented qualification exercise for one actual application. Candidate mixtures would first be screened at the required strength and age, followed by tests or accepted evidence covering the project's other governing conditions. The resulting eligible set could be assessed using common inventory conventions and explicit delivery and construction scenarios. This design would be more resource-intensive than reading public declarations, but it would answer a question the present study deliberately leaves open: how much of the observed declared dispersion survives full qualification? The value would come from linking documentary and engineering evidence, not merely enlarging the number of rows. Even a modest number of carefully matched alternatives could be more informative for that question than thousands of incompletely described products.

Public disclosure could also be improved through a compact, structured performance record associated with each mixture identifier. Such a record would distinguish specified from measured strength, state the reference age and relevant specimen convention, and identify density, exposure, air, and construction requirements where applicable. Adding fields creates a reporting burden and may encounter commercially

sensitive information, so it should prioritize attributes that determine whether a comparison is valid rather than demand a complete proprietary recipe. A verification statement could distinguish explicitly reported values from inferred classifications and unavailable information. The objective would be to make environmental results easier to connect to a qualified use, allowing future analyses to form defensible matched sets without relying on assumptions hidden in mixture names.

Conclusions

The analysis demonstrates substantial variation in declared embodied carbon within comparable strength classifications, while showing why the meaning of that variation depends on the source. In the NRMCA model, the lowest scenarios are 36.07-39.18% below the modeled zero-SCM reference within the six normal-weight strength classes. The separate DuPont case shows considerable dispersion among some products sharing a 28-day strength label. However, its missing mixture-level performance attributes prevent translating those ranges directly into substitution savings. The combined finding is that strength matching provides a useful way to investigate carbon differences, but the level of equivalence established must remain part of the reported result.

The original sensitivity calculations clarify where the numerical evidence is most useful. Large differences against a high-impact reference coexist with near-ties among lower-impact scenarios; changing the baseline changes the claimed incremental improvement; and dividing GWP by strength can make an index fall while per-volume emissions rise. None of these observations requires an invented statistical model. They follow from transparent arithmetic on the declared values, supplemented by an explicit account of the comparison's limits. Carbon headroom and ranking-reversal thresholds help direct further investigation, but they remain deterministic decision aids rather than estimates of uncertainty or whole-life benefit.

The resulting recommendation is to establish accounting compatibility, qualify the alternatives for the intended use, and then compare their environmental burdens at an explicit quantity and life-cycle boundary. Public EPDs can support that process by identifying promising candidates and documenting product-stage impacts. They cannot independently settle questions about structural redesign, service life, market-wide mitigation, or universal binder superiority. A valuable carbon comparison therefore ends with a traceable conditional conclusion: for the alternatives shown to meet the same requirements, under compatible assessment rules and stated quantities, the reported difference is relevant to the decision. Further claims require the additional evidence appropriate to their scale.

References

Anderson, J., & Moncaster, A. (2020). Embodied carbon of concrete in buildings, Part 1: Analysis of published EPD. *Buildings and Cities, 1*(1), 198-217. https://doi.org/10.5334/bc.59

Azarijafari, H. (2020, August). *Towards comparable environmental product declarations of concrete: Insights from a meta-analysis and probabilistic comparative LCA approach* (ACI CRC 18.516, Final report). ACI Foundation. https://www.acifoundation.org/Portals/12/Files/PDFs/Report_CRC18.516_0820.pdf

AzariJafari, H., Guest, G., Kirchain, R., Gregory, J., & Amor, B. (2021). Towards comparable environmental product declarations of construction materials: Insights from a probabilistic comparative LCA approach. *Building and Environment, 190*, Article 107542. https://doi.org/10.1016/j.buildenv.2020.107542

AzariJafari, H., Rangelov, M., Gregory, J., & Kirchain, R. (2023). Suitability of EPDs for supporting life cycle and comparative analysis of concrete mixtures. *Environmental Science & Technology, 57*(19), 7321-7327. https://doi.org/10.1021/acs.est.2c05697

Broyles, J. M., Gevaudan, J. P., & Brown, N. C. (2024). A compiled dataset of ready-mix concrete environmental product declarations for life cycle assessment. *Data in Brief, 52*, Article 109852. https://doi.org/10.1016/j.dib.2023.109852

CalPortland Company. (2019). *EPD for concrete produced at 9 CalPortland Washington facilities* (Environmental product declaration No. NRMCAEPD: 20026). National Ready Mixed Concrete Association. https://www.nrmca.org/wp-content/uploads/2019/10/CalportlandWA_EPD2019.pdf

Chen, C., Habert, G., Bouzidi, Y., Jullien, A., & Ventura, A. (2010). LCA allocation procedure used as an incitative method for waste recycling: An application to mineral additions in concrete. *Resources, Conservation and Recycling, 54*(12), 1231-1240. https://doi.org/10.1016/j.resconrec.2010.04.001

Damineli, B. L., Kemeid, F. M., Aguiar, P. S., & John, V. M. (2010). Measuring the eco-efficiency of cement use. *Cement and Concrete Composites, 32*(8), 555-562. https://doi.org/10.1016/j.cemconcomp.2010.07.009

Habert, G., Arribe, D., Dehove, T., Espinasse, L., & Le Roy, R. (2012). Reducing environmental impact by increasing the strength of concrete: Quantification of the improvement to concrete bridges. *Journal of Cleaner Production, 35*, 250-262. https://doi.org/10.1016/j.jclepro.2012.05.028

Habert, G., Miller, S. A., John, V. M., Provis, J. L., Favier, A., Horvath, A., & Scrivener, K. L. (2020). Environmental impacts and decarbonization strategies in the cement and concrete industries. *Nature Reviews Earth & Environment, 1*(11), 559-573. https://doi.org/10.1038/s43017-020-0093-3

International Organization for Standardization. (2006). *Environmental labels and declarations-Type III environmental declarations-Principles and procedures* (ISO Standard No. 14025:2006). https://www.iso.org/standard/38131.html

International Organization for Standardization. (2017). *Sustainability in buildings and civil engineering works-Core rules for environmental product declarations of construction products and services* (ISO Standard No. 21930:2017). https://www.iso.org/standard/61694.html

Juenger, M. C. G., Snellings, R., & Bernal, S. A. (2019). Supplementary cementitious materials: New sources, characterization, and performance insights. *Cement and Concrete Research, 122*, 257-273. https://doi.org/10.1016/j.cemconres.2019.05.008

Konradsen, F., Hansen, K. S. H., Ghose, A., & Pizzol, M. (2024). Same product, different score: How methodological differences affect EPD results. *The International Journal of Life Cycle Assessment, 29*(2), 291-307. https://doi.org/10.1007/s11367-023-02246-x

Miller, S. A. (2020). The role of cement service-life on the efficient use of resources. *Environmental Research Letters, 15*(2), Article 024004. https://doi.org/10.1088/1748-9326/ab639d

Miller, S. A., & Myers, R. J. (2020). Environmental impacts of alternative cement binders. *Environmental Science & Technology, 54*(2), 677-686. https://doi.org/10.1021/acs.est.9b05550

National Ready Mixed Concrete Association. (2022). *NRMCA member industry-average EPD for ready mixed concrete* (Environmental product declaration No. EPD10294). NSF Certification, LLC. https://www.nrmca.org/wp-content/uploads/2022/02/NRMCA_EPDV3-2_20220223.pdf

Olsson, J. A., Miller, S. A., & Kneifel, J. D. (2024). A review of current practice for life cycle assessment of cement and concrete. *Resources, Conservation and Recycling, 206*, Article 107619. https://doi.org/10.1016/j.resconrec.2024.107619

Provis, J. L. (2018). Alkali-activated materials. *Cement and Concrete Research, 114*, 40-48. https://doi.org/10.1016/j.cemconres.2017.02.009

Scrivener, K., Martirena, F., Bishnoi, S., & Maity, S. (2018). Calcined clay limestone cements (LC3). *Cement and Concrete Research, 114*, 49-56. https://doi.org/10.1016/j.cemconres.2017.08.017

UN Environment, Scrivener, K. L., John, V. M., & Gartner, E. M. (2018). Eco-efficient cements: Potential economically viable solutions for a low-CO2 cement-based materials industry. *Cement and Concrete Research, 114*, 2-26. https://doi.org/10.1016/j.cemconres.2018.03.015

U.S. Environmental Protection Agency. (2022, December 22). *Interim determination on low carbon materials under Inflation Reduction Act Sections 60503 and 60506.* https://www.epa.gov/system/files/documents/2023-01/2022.12.22%20Interim%20Determination%20on%20Low%20Carbon%20Materials%20under%20IRA%2060503%20and%2060506_508.pdf

Van Roijen, E., Sethares, K., Kendall, A., & Miller, S. A. (2024). The climate benefits from cement carbonation are being overestimated. *Nature Communications, 15*(1), Article 4848. https://doi.org/10.1038/s41467-024-48965-z

Xi, F., Davis, S. J., Ciais, P., Crawford-Brown, D., Guan, D., Pade, C., Shi, T., Syddall, M., Lv, J., Ji, L., Bing, L., Wang, J., Wei, W., Yang, K.-H., Lagerblad, B., Galan, I., Andrade, C., Zhang, Y., & Liu, Z. (2016). Substantial global carbon uptake by cement carbonation. *Nature Geoscience, 9*(12), 880-883. https://doi.org/10.1038/ngeo2840