# Depth Before Derivation: Meta-Level Modification in Orthographic and Morphological Stabilization

Flyxion

August 2026

**Abstract**

*Embodied Constraint and Minimal Assembly* models the stabilization of the Arabic script as a cost–distinctness optimization: graphemes persist when a proposed modification's informational gain exceeds its production cost. The model is powerful but flat. It treats every stabilizing event in the script's history as an instance of the same operation: a single-level trade between $C$ and $D$ applied directly to a grapheme. This essay argues that the script's own history contains at least two operations of different kind, not merely different degree. Dotting a letter and regularizing a script's joining rules are not the same type of event. The first modifies a grapheme. The second modifies the rule that generates a grapheme's possible forms. Read through an admissibility-dynamics framework, in which a modificatory operation is classified by its meta-depth $k$, the distinction becomes visible as a structural fact about the script rather than a stylistic observation about it. The same distinction reappears in Arabic derivational morphology, where the traditional measures built on the paradigm root fa'ala show the pattern with unusual clarity: applying a known measure to a root is depth-0, while the historical fixing of the measure set itself is depth-1. This reframing does not replace the cost–distinctness model; it adds an axis the model does not have, and the resulting picture predicts which parts of a script's, or a morphology's, history should be irreversible in a different sense than others.

## 1 A Flat Model of a Layered History

*Embodied Constraint* treats the stabilization of a writing system as an equilibrium condition. A grapheme $g$ persists in a stable configuration when, for every available augmentation $a$, the ratio of informational gain to production cost fails to justify adopting $a$:

$$S = \alpha D - \beta C,$$

with the script settling near local maxima of $S$. Section 4 of that essay applies this model to dotting: a minimal graphic addition to a skeletal base that yields a categorical phonemic distinction, at very low additional cost. Section 14 applies structurally the same model to a different phenomenon: the historical replacement of unbounded, context-sensitive ligature forms with a fixed four-way positional system (isolated, initial, medial, final). Both are described as instances of the same underlying dynamic, compression under constraint, cost against distinctness, equilibrium found and held.

The description is not wrong. But it collapses two operations that differ in what they act *on*. Dotting changes a grapheme. Positional regularization changes the rule that assigns a grapheme its permissible shapes. The cost–distinctness formalism, as stated, has no way to mark this difference; $C(a)$ and $D(g_1, g_2)$ are defined at the level of individual graphemes throughout, and the model offers no separate variable for modifications that operate one level up, on the generative rule itself.

## 2 Two Kinds of Modification

Consider the two cases side by side.

**Dotting.** Before systematic dotting, dāl and dhāl, د and ذ (to use a rough transliteration) shared an undotted skeletal base $b$. Dotting introduces an augmentation $a$ such that $g_1 = b$, $g_2 = a(b)$. The operation acts directly on the space of admissible graphemes: it adds one new admissible form. Everything else about the writing system, its joining rules, its positional behavior, its inventory of augmentation types, is untouched. The operation reads the current state of the grapheme inventory and writes one new member into it.

**Positional regularization.** Before four-position stabilization, the shape a letter took depended on its specific left and right neighbors, $F(g_i \mid g_{i-1}, g_{i+1})$, a function whose domain grows combinatorially in the number of graphemes. Regularization does not add a new grapheme, and it does not disambiguate two previously conflated forms. It replaces the function $F$ itself with a lower-complexity function $F(g_i \mid p)$, where $p$ ranges over four positions rather than over all neighbor pairs. The object being modified is not a grapheme. It is the rule that had been generating grapheme forms.

These are not two points on the same axis of "how much change occurred." They are modifications with different targets. One writes into the space of admissible forms under a fixed generative rule. The other rewrites the generative rule.

## 3 A Depth Axis for Orthographic Modification

This distinction has a name in admissibility-dynamics work developed elsewhere in this program: meta-depth. A modificatory operation is depth-0, written $M^{(0)} : R \to R$, if it acts directly on the admissibility structure $R$, the set of currently reachable or permitted configurations. An operation is depth-$k$ for $k \geq 1$, written $M^{(k)} : M^{(k-1)} \to M^{(k-1)}$, if it acts on the machinery that itself modifies $R$, rather than on $R$ directly. The depth axis is inductive: there is no need for an $R^{(-1)}$ object, only a base case and a recursive step upward.

Under this axis, dotting is depth-0. It is a write operation on the admissible-grapheme set $R_t$ at a given moment: one new distinguishable form is added, everything about how future forms will be generated is left alone. Positional regularization is depth-1. It is a write operation not on $R_t$ but on the rule $F$ that determines what $R_{t+1}$ can look like given $R_t$ and a neighboring context; it changes the generative machinery, not a product of that machinery.

This is not a difference of magnitude. A very large dotting campaign, applied to every skeletal ambiguity in the alphabet at once, would still be depth-0: many individual writes to $R_t$, executed together, but each one a disambiguation of an existing grapheme pair under an unchanged rule for how graphemes get their shapes in context. Positional regularization, even applied to a single letter family first and generalized later, is depth-1 from the moment it changes how shape is determined at all, because after that point every subsequent grapheme in the affected family is generated by a different function than before.

## 4 The Measures of *f*-ʿ-*l*: Templates as Bounded Generators

The same depth distinction reappears, more sharply, in Arabic morphology, where it can be shown rather than only argued for, because the tradition itself already names its generative apparatus with a fixed exemplar root. Arabic grammar traditionally illustrates its derivational templates, the measures or awzān, using the root *f*-ʿ-*l* (faʿala, فَعَلَ), whose three consonants stand in generically

for $C_1$, $C_2$, $C_3$ in any root. The measures are conventionally cited by name using this root directly: Form I is called *fa‘ala*, Form II *fa“ala*, and so on, the pattern name and the paradigm word being the same string. This is already a small piece of evidence for the argument to follow: the tradition treats the measures as a closed, nameable inventory precisely because they function as a fixed generative rule, not as an open field of possible modification.

Applying fa‘ala’s own consonants to the ten canonical measures gives:

| Form | Pattern | Arabic | Gloss (root sense: to do/act) |
|---|---|---|---|
| I | fa‘ala | فَعَلَ | he did |
| II | fa“ala | فَعَّلَ | he made [someone] do; he did intensively |
| III | fā‘ala | فَاعَلَ | he did [something] with/against [someone] |
| IV | ’af‘ala | أَفْعَلَ | he caused to do |
| V | tafa“ala | تَفَعَّلَ | he had himself made to do; did gradually |
| VI | tafā‘ala | تَفَاعَلَ | he pretended to do; they did to one another |
| VII | infa‘ala | اِنْفَعَلَ | it got done (passive/inchoative) |
| VIII | ifta‘ala | اِفْتَعَلَ | he did for himself |
| IX | if‘alla | اِفْعَلَّ | he became [a quality, typically color/defect] |
| X | istaf‘ala | اِسْتَفْعَلَ | he sought to have [it] done |

Every one of these forms is generated from the same three consonants, $(C_1, C_2, C_3) = (f, \text{`}, l)$, by a different composite operator drawn from a small set of primitive moves: vowel insertion, gemination of $C_2$ (Form II), vowel lengthening after $C_1$ (Form III), prefixation of hamza (Form IV) or *ta-* (Forms V, VI) or *ist-* (Form X), infixation of *-t-* after $C_1$ (Form VIII), or a nasal prefix with final gemination (Form VII, IX). In the notation of Appendix D of *Embodied Constraint*, each measure is a composite $T = A_k \circ \cdots \circ A_2 \circ A_1$ built from these primitives, and applying a known measure to a root is a single write: it inserts one fully specified word $W$ into the space of forms derivable from $(f, \text{`}, l)$, using a generative rule, the finite measure set $\mathcal{T} = \{T_{\mathrm{I}}, \ldots, T_{\mathrm{X}}\}$, that the operation itself leaves untouched. This is the morphological analogue of dotting. It is depth-0: a write to $R_t$, the space of admissible derived words for a given root, under a fixed rule for how derivations happen.

The measure set itself is the depth-1 object. Ten (extended, in some grammars, to fifteen for rarer patterns) canonical templates is a closed inventory imposed on what could, in principle, be an unbounded combinatorial space of primitive-operator compositions: any sequence of vowel insertions, geminations, prefixations, and infixations applied to a triliteral root defines some derivable word, but only a small, named, pedagogically fixed subset of these compositions counts as a productive measure that a speaker is expected to recognize and extend to new roots. The historical stabilization of that inventory, whatever process fixed it at approximately ten forms rather than leaving derivation as an open field of ad hoc consonant-vowel recombination, is structurally the same event as positional regularization in the script: a reduction of an unbounded generative space, $F(g_i \mid g_{i-1}, g_{i+1})$ in one case, arbitrary compositions of morphological primitives in the other, to a small, closed, teachable rule. Learning a new root and applying Form II to it is depth-0. Recognizing or proposing an eleventh productive measure, one that would apply generally across roots rather than existing as an isolated lexical fossil, would be depth-1, and this is in fact why such proposals are rare and contested in the grammatical tradition: a depth-1 change is not merely a bigger version of a depth-0 one, it is a change to what counts as a legitimate way of generating a word at all.

This also sharpens a point *Embodied Constraint* makes about predictive generativity (Section 11.4) without quite stating it in these terms. The essay observes that speakers can infer the sense of an unattested derived form by applying a known measure to an unfamiliar root. That predictive capacity is only available because measure application is depth-0: it presumes a fixed, already-learned generative rule and simply executes it on new input. If measure formation were itself a live, per-word option, generalization from known cases to new ones would not be reliable in the way the essay describes, since there would be no stable rule to apply. The productivity of Arabic derivational morphology is, in this sense, a direct consequence of keeping template formation at depth-1, rare and historically settled, while leaving template application at depth-0, frequent and freely generalizable.

## 5 Gemination as a Fixed Mark: Shaddah and a Comparative Aside

Shaddah (شَدَّة), the diacritic marking consonant doubling, is depth-0 by the same test applied to dotting and sukūn: it writes a single explicit fact, "this consonant is geminate here," into $R_t$, the space of admissible syllabic readings, without touching any rule for how future marks of this kind get assigned. It is, in the vocabulary of Appendix D, the operator $G(C) = CC$: identity duplication rather than the introduction of a new graphic primitive, cheaper in $C(a)$ than dotting for exactly that reason, since no new shape has to be learned, only a repetition of one already known.

A comparative case outside Arabic sharpens what depth-0 status does and does not guarantee about a mark's later history. Spanish ñ descends from Latin geminate *nn*, as in *annu* → *año*. Medieval scribes abbreviated a doubled letter by writing one instance of it with a small superscript wave standing in for the second, a purely economical compression, the same motive Appendix B attributes to ligature simplification in Arabic: fewer strokes for the same string. At that stage the tilde was depth-0 in the weakest sense, a shorthand for an already-present doubling, carrying no phonemic information the doubled letters themselves did not already carry. Only later, once the geminate nasal had undergone palatalization into a distinct phoneme, the palatal nasal, did the same mark come to function the way a dot on dāl/dhāl functions: a categorical phonemic distinguisher, separating *año* from a hypothetical *ano* rather than merely abbreviating a spelling. The graphic object did not change. Its function migrated from compression to distinction without any change in depth, both stages being depth-0 writes, which shows that depth and function-type, compression versus differentiation, are independent axes. A mark can move from one function to the other while never leaving depth-0, so long as neither stage rewrites the rule generating the space of marks available to writers.

## 6 Person, Tense, and Mood: Composing Two Fixed Paradigms

The measures discussed above generate a derived stem from a root; they do not by themselves produce a finished, usable verb. A further layer, inflection for person, gender, number, tense-aspect, and mood, applies on top of the stem through a second, independently fixed paradigm of prefixes and suffixes. Holding the root and measure constant at Form I of *f*-ʿ-*l*, the imperfect indicative alone yields:

| Form | Arabic | Person/gender | Gloss |
|---|---|---|---|
| ’af‘alu | أَفْعَلُ | 1st singular | I do |
| naf‘alu | نَفْعَلُ | 1st plural | we do |
| taf‘alu | تَفْعَلُ | 2nd masc. singular / 3rd fem. singular | you do / she does |
| yaf‘alu | يَفْعَلُ | 3rd masc. singular | he does |
| yaf‘alūna | يَفْعَلُونَ | 3rd masc. plural | they do |

Each of these is generated by a fixed paradigm of person-marking affixes, $\mathcal{P}$, applied to the stem produced by the measure paradigm $\mathcal{T}$. Neither paradigm modifies the other. Selecting a measure and selecting a conjugational affix are independent depth-0 writes, each against its own closed rule set, and a fluent verb form is simply the composition of both: stem $= T(\text{root})$, word $= P(\text{stem})$. This composability is itself informative. It is only because $\mathcal{T}$ and $\mathcal{P}$ are each separately fixed at depth-1 that their composition at depth-0 remains predictable, a speaker who knows both paradigms can generate or parse a form never previously encountered.

Mood marking on the imperfect verb makes the connection back to sukūn explicit rather than analogical. The three moods of the Arabic imperfect are distinguished only by the final short vowel or its absence: indicative *yaf‘alu* (damma), subjunctive *yaf‘ala* (fatḥa), and jussive *yaf‘al* (يَفْعَلْ, no vowel at all). The jussive is formed by exactly the operation Lesson 10 of the Qaedah names Jazm: a sukūn placed on the final consonant, closing the word where an open syllable had stood. The traditional Arabic grammatical term for the jussive mood, *majzūm* (مَجْزُوم), is built on the same root as *jazm* (جَزْم) itself, the term is not merely analogous to sukūn, it is named for it. The same depth-0 operator that closes a mid-word syllable in ordinary orthography, marking absence of a following vowel within a fixed system that otherwise presumes vocalic continuation, reappears one grammatical level up as the mechanism that distinguishes an entire mood. Sukūn’s status as a null-state write, argued above to be a boundary case between marking presence and marking absence, turns out to do exactly the same job whether it closes a syllable inside a word or closes the final syllable of an inflected verb to signal jussive mood. The mark did not change. Only the grammatical weight riding on it did.

## 7 Waqf: Sukūn at the Boundary of the Utterance

The pattern identified for the jussive mood, sukūn’s closure operator reused one level up from the syllable to close an entire grammatical category, recurs a second time at a still larger scope. Arabic nouns in continuous recitation or connected reading carry case-and-definiteness endings: an indefinite noun in the nominative takes tanwīn, the doubled ḍamma pronounced as *-un*, so that the bare consonantal name of the letter alif, cited indefinitely, is read *alifun* (أَلِفٌ), “an alif,” against the definite *al-alifu* (الْأَلِفُ), “the alif,” where the definite article replaces the indefinite marking rather than combining with it. This ending is not optional ornamentation; the presence or absence of tanwīn is what signals indefiniteness in the first place, doing morphologically what the English article *a* does syntactically.

At a pause, whether the pause is dictated by the sense of a sentence, by Qur’ānic waqf convention, or simply by a reader stopping for breath, this ending is dropped. *Alifun* at pause is read *alif*, its final vowel and nunation removed regardless of whether the underlying word was indefinite, definite, nominative, or genitive; the rule governing this is stated in the Qaedah itself in exactly sukūn’s own terms, that the last letter of a paused word should be made *sakin*. The operation is not

a new mechanism invented for pause. It is the identical closure operator examined in Section 10.2 of *Embodied Constraint*, applied at the largest scope encountered so far: not a syllable within a word, not a mood-marking morpheme on a verb, but the boundary of an entire phrase or utterance. A reader moving through unvocalized or lightly vocalized continuous text, a newspaper column being the practical case, uses this drop precisely as a period or comma is used in scripts that mark punctuation graphically: the disappearance of the expected final sound is itself the signal that a unit has closed, doing prosodically what a mark on the page does visually in other traditions.

Three scopes, one operator. Word-internal sukūn closes a syllable inside a still-continuing word. Jussive sukūn closes a verb's mood category. Waqf closes a phrase. In every case the operation is the same depth-0 write against the same background assumption, that vocalic continuation is the default and its absence must be explicitly marked, and in every case the mark's scope, not its identity, is what changes. This is the clearest demonstration in the corpus examined here that depth and scope are themselves independent: an operation can remain at depth-0, a direct write to $R_t$ with no modification to the rule generating $R_t$, while the size of the unit it closes grows from a single consonant to an entire spoken sentence.

## 8 Hamza and the Selection of a Seat

Hamza, the glottal stop, is written on one of several carriers depending on its surrounding vowel environment: on alif, on wāw, on yā', or unsupported on the line, chosen by a rule that weighs the strength of the vowel before and after the hamza against a fixed ranking, kasra strongest, ḍamma next, fatḥa weakest, sukūn weakest of all, and selects the carrier matching the stronger of the two. Writing *su'āl* with hamza on wāw, or *ri'āsa* with hamza on yā', is an application of this ranking to a specific word, a depth-0 write choosing one seat from among the admissible options given the surrounding vowels, executed against a fixed comparative rule the operation itself does not touch. *Embodied Constraint* notes, in Section 10.4, that hamza functions as a boundary marker between silence and articulated sound. That observation concerns hamza's phonetic role. The seat-selection rule is a separate, orthographic fact layered on top of it, and it is this second layer that shows the depth structure: early Qur'ānic manuscripts are known to have treated hamza inconsistently, at times omitting it from the rasm entirely, before grammatical codification fixed both its notation and its seat-selection ranking into the closed algorithm later scribes and readers inherited as given. The codification of the ranking was depth-1. Every subsequent application of it, one hamza, one word, one carrier chosen, is depth-0.

## 9 Noon Sākinah and Tanwīn: A Second Partition of the Consonants

Section 5 of *Embodied Constraint* reads the sun-letter and moon-letter division as a phonologically motivated partition of the consonant inventory, coronal versus non-coronal, governing whether the lateral of the definite article assimilates. A structurally similar but independently motivated partition governs what happens when noon sākinah or tanwīn is followed by a consonant. Four treatments are recognized. Idhār, clear pronunciation with no modification, applies before the throat letters, the same huroof-e-halqi that block ghunnah in Lesson 5 of the Qaedah. Idghām, assimilation of the noon into the following consonant, applies before a small class dominated by nasals and liquids, lām, rā', nūn, mīm, wāw, and yā' (ل ر ن م و ي), with ghunnah retained for four of these and dropped for two. Iqlāb, conversion of the noon sound itself into a meem before articulation, applies solely before bā' (ب). Ikhfā, a concealed, partially nasalized pronunciation,

applies before the remaining consonants, the majority of the inventory. This four-way classification is a second closed taxonomy imposed on the same twenty-nine consonants the sun-letter division already partitions, keyed this time not to coronality alone but to a mix of articulatory distance from noon's own place and manner, nasals and liquids cluster toward assimilation, the labial bā' (ب) alone earns full conversion rather than mere assimilation, sharing noon's nasal manner once converted to meem, and the throat letters, most distant from noon's place of articulation, block modification entirely. As with the sun letters, applying the correct treatment to a specific instance of noon sākinah meeting a specific following consonant is depth-0, a lookup against a fixed four-way rule. The historical settling of which consonants belong to which of the four classes, a classification a beginning reader must simply memorize rather than derive, is the depth-1 fact making the depth-0 lookups possible at all. That two independently motivated closed partitions of the same twenty-nine-letter inventory, one for the definite article, one for noon sākinah, coexist without needing to agree with one another is itself evidence that depth-1 stabilization in this system happens locally, rule by rule, rather than as a single global classification of the consonants that every subsequent process must respect.

## 10 Broken Plurals: A Case the Depth Axis Does Not Close Cleanly

Not every derivational domain in Arabic resolves as cleanly into a depth-0/depth-1 pair as the verb measures do. Plural formation offers a useful complication. Some nouns pluralize by suffixation alone, a sound plural formed with *-ūn*/*-īn* for masculine or *-āt* for feminine, a straightforward depth-0 affixation against one of two fixed suffix paradigms, structurally identical to the person-marking affixes discussed above. The majority of nouns, however, pluralize by internal vowel-pattern change, a broken plural, *jam‘ taksīr*: *kitāb* becomes *kutub*, *qalam* becomes *aqlām*, *rajul* becomes *rijāl*. This looks, at first, like the nominal counterpart of the verbal measures examined above, a fixed set of templates applied to a triliteral base. But the broken-plural template set is neither small nor fully productive in the way the ten verbal measures are. Traditional grammars enumerate roughly thirty common patterns, and critically, which pattern a given singular noun takes is only partially predictable from its phonological shape; a substantial number of broken plurals must be learned per noun rather than derived from a rule a speaker could apply to an unfamiliar root with confidence. This resists a clean depth assignment. It is not depth-0 in the sense measure-application to fa‘ala is depth-0, since depth-0 there presumes a determinate, executable rule, and no such rule reliably determines which of thirty patterns a new noun will take. Nor is it depth-1, since no individual instance of forming a broken plural modifies the pattern inventory itself. The most accurate description borrows the interrogative/modificatory distinction developed elsewhere in this program's admissibility-dynamics work: retrieving the correct broken-plural pattern for a specific noun functions less like the execution of a generative rule and more like an interrogative lookup against a large, only loosely rule-governed stored table, a read operation over accumulated lexical fact rather than a write executed against a closed formal grammar. The verb measures earn their productivity, in Section 11.4 of *Embodied Constraint* and the discussion above, from sitting cleanly at depth-0 under a closed depth-1 rule. Broken plurals show what a derivational domain looks like when that condition only partially holds, and the honest response is not to force a depth number onto it but to note that not all of Arabic morphology is measure-like, some of it is closer to memorized association than computed derivation, however much a given broken plural pattern resembles a measure once it is already known.

## 11 What the Cost–Distinctness Model Cannot See

The consequence of collapsing these into one axis is that the cost–distinctness model has no principled account of why some orthographic changes are effectively irreversible in a stronger sense than others. Under $S = \alpha D - \beta C$, an unwound dotting decision (returning to an undifferentiated skeletal form) is symmetric with adopting one: both are single moves along the same $D$-$C$ tradeoff, evaluated at the grapheme level, and either could in principle be favored again if the weighting parameters shifted. But undoing positional regularization is not symmetric with adopting it in the same way. Reverting from four-position allographs back to fully context-sensitive ligature forms would not just cost more; it would require re-deriving, or re-inventing, the entire combinatorial mapping $F(g_i \mid g_{i-1}, g_{i+1})$ that regularization eliminated. The cost of that reversal is not a point on the same $C$ scale as the cost of undotting a letter, because it is a cost incurred at the level of the rule generating the grapheme space, not at the level of a grapheme.

Put in terms of the essay's own resource-cost language: a depth-0 reversal pays $C(a)$ for a single grapheme. A depth-1 reversal pays for reconstructing an entire generative function, which is not $C(a)$ summed over graphemes but a qualitatively different kind of cost, closer to what the admissibility framework treats as the cost of restoring lost machinery rather than the cost of restoring a lost state. *Embodied Constraint*'s own appendices gesture toward this without naming it: Appendix B's treatment of ligature suppression as "entropy reduction" correctly identifies that something more than grapheme-level compression is happening, that the domain of the joining function itself is shrinking from $\mathcal{O}(n^2)$ to $\mathcal{O}(4n)$, but the essay does not connect this observation back to its own $S = \alpha D - \beta C$ equilibrium model, because that model has no slot for a change in the generating rule as distinct from a change in the generated output.

## 12 Sukūn as a Boundary Case

One diacritic in the same corpus resists clean classification under either label, and the resistance is informative. Sukūn does not add a new admissible grapheme, the way dotting does, and it does not rewrite a generative rule, the way positional regularization does. It marks the absence of a following vowel: a null state made explicit within a system that otherwise treats vowel continuation as a default. In the vocabulary introduced above, sukūn is a depth-0 operation, since it acts directly on $R_t$, the space of admissible syllabic readings for a given consonant, but its content is unusual for a depth-0 write: rather than adding a distinguishable positive form, it declares that no further form follows. This is closer to what admissibility-dynamics work elsewhere calls a constitutive operation whose payload is a closure condition rather than a new member. It is worth flagging as a case that a two-term cost–distinctness model handles by accident (its $C$ is low, its $D$ from an unmarked alternative is high, so it fits the same inequality as dotting) without the model registering that it is doing something categorically different from adding a new letter.

## 13 Conclusion

*Embodied Constraint*'s central claim, that scripts stabilize under recurring pressure toward informational distinctness at bounded cost, does not need to be abandoned to notice that the pressure operates on at least two different objects across the script's documented history: the grapheme inventory itself, and the rule that generates grapheme forms from context. Treating both under a single $S = \alpha D - \beta C$ equilibrium obscures a distinction that the script's own evolution makes visible when the objects of modification are tracked rather than assumed uniform. Dotting operates within

a fixed rule for what a letter can look like. Positional regularization operates on that rule. The difference is not one of scale along a shared axis; it is a difference in what kind of thing was written to. A model of orthographic stabilization that wants to explain why some historical changes to a script are recoverable in a way others are not needs a depth axis, not only a cost axis.