Rendering Mutual Fund Risk/Return Summary Taxonomy 2010 Instances for the SEC Viewer

Date 2010-04-06

1 Goals

The purpose of this document is to explain new layout rules applied by the Commission Viewer's "rendering engine", that is, how the engine transforms XBRL instances into HTML output for display in a web browser, with a particular focus on instances of the Mutual Fund Risk/Return Summary Taxonomy 2010.

For instances of the Mutual Fund Risk/Return Summary Taxonomy 2010, it is important to be able to embed tabular displays into text blocks for rendering. The approach is general enough to apply to Interactive Data filings that are valid with respect to EDGAR Filer Manual Volume II chapter 6, version 13.

Layout means the overall structure of a report: rows, columns, heading rows, heading columns, and nesting of reports within others. There is almost no change to the Formatting, such as how a value is shown in a cell, such as 0 being represented by the word "none", dash, "0", etc., with or without scaling, as HTML or as plain text.

Table of Contents

2 Terminology

The audience of this document is assumed to have some familiarity with XBRL technical terms; as used in this document the following XBRL terms figure prominently.

1. **Fact:** A number, date, fragment of plain text, or formatted section of a document that has been "tagged" with an XBRL element and assigned a *Context*.
2. **Context:** The properties that a set of facts has in common. A context includes the period to which the fact applies, its unit of measure (for numeric facts), and may include the document in which it is shown (typical only for textual facts), what Fund Series it applies to (if any), what share class it applies to (if any), and, if it is a performance measure, other properties such as whether it is before or after tax, or whether it is an external benchmark.
3. **Axis and Domain Member** elements: these represent properties other than period and unit that appear in the context, and the values they take on. For example, "Fund Series" is an axis and "The Cabbage Fund" and "The Eggplant Fund" are domain members of that axis. "Share Class" is an axis and "Eggplant Fund Class C" is a domain member of that axis.
4. **Primary** element: Each fact is an occurrence of a named concept such as "Management Fees" or "Risk Narrative"; the primary element is the one that is usually (though not always) shown in the labels that run down the left side of a table.
5. **Presentation Group:** A group of one or more hierarchies of primary elements, and of axes and their domain members. Technically, it is an XBRL "base set".
6. **Presentation Hierarchy:** An arrangement (set of relationships) of primary or axis and domain elements that informs the rendering engine how to filter, to nest, order, facts that match its primary elements or members. The result is conventionally a table. The presentation hierarchy furthermore informs the engine how to label the row and column headings, and (in some cases) how to display facts.
7. **Report:** When facts in an instance are displayed using the hierarchies in a Presentation Group, the result is called a Report.
8. **Presentation Linkbase:** A presentation linkbase is formally a file containing a set of presentation groups but the file itself is unimportant because it usually means "all the information linked to an instance that may be used to render that instance", so, for example, ignoring any relationships that have been prohibited or overridden by taxonomy extensions.
9. **Preferred Label**: Whenever a column or row heading is displayed, the presentation linkbase may select among several text labels to display; the label selected is the preferred label.

3 Motivation

A presentation hierarchy of primary elements is generally adequate for conventional income statements, balance sheets, and cash flow statements. In these displays, the primary items always appear on the leftmost column, and the columns always contain the data for a specified period.

To display the tabular displays required by the Risk/Return Summary as specified on SEC Form N-1A, it is necessary to:

1. Order the members of an axis so that facts in contexts that specify different members of the axis are displayed in a desired order. For example, expense data about share classes A, B, C, I and R would normally be displayed in that (alphabetical) order.
2. Nest the axes, so that, for example, a table could show one set of share classes for a series (e.g., Equity fund classes A and B) followed by other classes of that series (Equity fund classes C and D); or vice versa (Classes, then Series).
3. Pivot the axes, so that some axes that by default would appear as column headings are moved to be groupings of row headings; or vice versa.
4. Select the contexts whose facts will be displayed in a table (for example, show only facts with the Series axis and member “Eggplant Fund”).

Also, a presentation hierarchy limited to primary items may be adequate when the hierarchy contains entirely numeric or date concepts, or when it consists entirely of text blocks. But to display the Risk/Return Summary as specified by Form N-1A while still displaying the interactive data facts that happen to be embedded in text it is also necessary to:

1. Interleave text narratives and headings that flow across the entire width of a logical page, with tables and images, all together on a single logical page.
2. Create a single Interactive Data document in which several distinct prospectuses for different groupings of fund series, and subsets of classes, on a single page.
3. Create distinct styles of Interactive Data views of the same facts, some that emphasize narrative and layout, and some that emphasize the presence of each individual fact in a “data dump” style of view.

The purpose of “embedding commands” is to give the preparer an improved ability to filter and organize data into a table that approximates a conventional, familiar layout, and also to position that table in place of where a text block displaying the table would have gone instead.

4 Specializing presentation groups

A conventional rendering of Interactive Data instances treats all relationship groups equivalently. The presentation group orders and indents the primary items in the vertical (rows) direction; the columns are generally for contexts, and therefore, for different units, periods and segments. As an exception, the SEC Viewer relies on patterns of words appearing in the definition text of a relationship group to provide special handling, as documented at http://www.sec.gov/spotlight/xbrl/staff-interps.shtml

> 1. Facts are arranged into columns with descending periods from left to right without regard to anything in the taxonomy.
> 2. Facts are shown in the same column if their end date and instant are the same, again without regard to anything in the taxonomy.

3. Facts are then shifted one column to the left if there is a "period start" preferred label, but only if the role definition contains the phrase "cash" followed by "flow".
4. Rows are grouped together into alphanumerically ordered groups, according to the contents of the segment elements (and scenario elements, although scenario elements are rejected by EDGAR).
5. If the role definition contains these word combinations, then the periods become rows and the Equity Class axis are the columns: (a) “stockholder” and “equity” but not “parenthetical” (b) “shareholder” and “equity” but not “parenthetical” (c) “partners” and “capital” but not “parenthetical” (d) “changes” and “equity” but not “parenthetical”.

Furthermore, if the role definition matches the pattern “sort - token - text” then only the text is shown in the SEC Viewer. For example, “0002 - Disclosure - Shareholder Fees” renders as simply “Shareholder Fess”.

These general layout rules are not changing, for backward compatibility. However, there are now additional ways of making overall changes to a layout by modifying the description text. When the tokens “Unlabeled”, “Elements”, or “Transposed” appear in the definition text between curly braces “{…}”, then one of the three transformations apply.

4.1 {Unlabeled}

When the definition text contains the token {Unlabeled}, then the elements’ labels do not appear in the leftmost column. Suppose, for example, there is a presentation group with the description “0000 - Document - Summary of Objectives”. Then a tabular display of the instance might look like this, with XHTML formatting of the text blocks providing italicization:

Summary of Objectives	
Objectives [Heading]	Investment Objectives
Objective, Primary [Text Block]	The primary objective of the Cabbage Fund is *long term capital appreciation*.
Objective, Secondary [Text Block]	A secondary objective is *preservation of capital*.

With description “0000 - Document - Summary of Objectives {Unlabeled}” the resulting display would have looked like this instead:

Investment Objectives
The primary objective of the Cabbage Fund is *long term capital appreciation*.
A secondary objective is *preservation of capital*.

The use of the {Unlabeled} token is not limited to presentation groups consisting solely of text and text blocks, but that is where it is most useful.

4.2 {Elements}

When the text contains the token {Elements}, then the display has an additional column to the right of the labels, showing the element id.

For example, consider a hypothetical example: a presentation group with description text “00 - Document - Example”:

Annual Fund Operating Expenses	
Management fees	0.66%
Other expenses	0.31%
Total annual fund operating expenses	0.97%

With the presentation group description “00 - Document - Example {Elements}” the elements’ names are shown alongside the label text given to them by the filer:

Shareholder Fees	Element	
Management fees	rr_ManagementFeesOverAssets	0.66%
Other expenses	rr_OtherExpensesOverAssets	0.31%
Total annual fund operating expenses	rr_ExpensesOverAssets	0.97%

If {Unlabeled} and {Elements} tokens both appear in the description, then {Elements} takes precedence.

An {Elements} display is sometimes called a “details group” or “data dump” because it is useful for reviewing all the detailed content of an Interactive Data instance.

4.3 {Transposed}

When the text contains the token {Transposed}, then the entire display, after any other layout has been applied, is transposed. That puts the rows on the columns, and vice versa. For example, with the presentation group “01 - Document - Example” a tabular display might look like this (assuming for illustration only that different years’ fees were actually shown in the same table):

Annual Fund Operating Expenses	2009	2008
Management fees	0.66%	0.66%
Other expenses	0.31%	0.29%
Total annual fund operating expenses	0.97%	0.95%

With the presentation group “01 - Document - Example {Transposed}” the same data would be displayed like this instead:

	Management fees	Other expenses	Total annual fund operating expenses
2008	0.66%	0.29%	0.95%
2009	0.66%	0.31%	0.97%

If the {Transposed} token appears with token {Elements} or token {Unlabeled} then {Transposed} takes precedence.

Badly formed syntax such as {Transposed,Elements} or {} is ignored.

Note that periods (2009 and 2008 in this example) sort descending when they are columns, but ascending when they are rows.

5 Filtering and sorting contexts

A conventional rendering of Interactive Data instances uses relationships in a presentation group to:

(a) Select (that is, filter) the facts to be shown on an Instance Report,

(b) Order (and indent) the line items,

(c) Indicate the preferred label to line items, and

(d) Indicate whether a numeric value should be shown with a “sign flip”.

Presentation groups now serve the first three of those four functions with respect to context axes and their domain members:

(a) Select (that is, filter) the contexts to be shown on a report,

(b) Order (but with only *one* level of indenting) domain member row and column headings, and

(c) Indicate the preferred labels to be used in row and columns for domain members.

5.1 Filtering

To determine whether a context is available for rendering in a presentation group, apply these criteria in this order:

1. If the context contains a member of an axis, and that axis element has presentation children in the group, but the member is not among its presentation children, then the context is not available.
2. If the presentation group contains an axis, and the axis has a default member, but that default member is not one of the presentation children of the axis, then the context is not available.
3. Otherwise, the context is available.

For example, suppose there is a presentation group in which the element dei_LegalEntityAxis has just one child eg_S0005977Member. The default member of dei_LegalEntityAxis is dei_LegalEntityDomain.

A context with the following segment will be unavailable, because it meets criterion 1: eg_S0001234Member is not in the presentation group:

```
<xbrli:segment><xbrldi:explicitMember
      dimension="dei:LegalEntityAxis"
      >eg:S0001234Member</xbrldi:explicitMember></xbrli:segment>
```

A context with the following segment will be unavailable because it meets criterion 2: the *default* member of the LegalEntityAxis is not in the presentation group (irrespective of whether dei:DocumentInformationDocumentAxis appears in the presentation group):

```
<xbrli:segment><xbrldi:explicitMember
      dimension="dei:DocumentInformationDocumentAxis"
      >eg:RetirementMember</xbrldi:explicitMember></xbrli:segment>
```

Informally: if a particular axis isn’t mentioned in a presentation group, then that axis won’t have any effect at all on which contexts are displayed. When an axis does appear, then the default member of an axis is usually in that presentation group so as to control how it is ordered with respect to other members.

Presentation groups have no impact on filtering of units or periods.

5.2 Context Ordering

In a conventional rendering of Interactive Data instances, facts are displayed in a table whenever they appear in the presentation group's hierarchy of primary concepts. Those facts are ordered top-to-bottom according to the hierarchy of primary concepts, and the facts are ordered left-to-right according (mainly) to their period. Other properties of the facts, such as the unit of measure (for example, USD vs. CAD), or properties of the context (for example, Legal Entity axis member) have a fixed precedence ordering that causes them to be displayed in columns and groupings that the presentation linkbase cannot influence.

Presentation groups now also control the order in which facts in an available context (available, in the sense of having passed filtering) will appear. Independently of whether an axis is displayed vertically (as a column) or horizontally (as a column), the order of an axis' children determines the order in which facts are displayed. For example, suppose there are facts in three contexts, all with the same period:

rr_ProspectusShareClassAxis – empty (default member rr_ShareClassDomain)
rr_ProspectusShareClassAxis – eg_C0005977Member
rr_ProspectusShareClassAxis – eg_C0006000Member

Assume this presentation group:

Element	**Order**
rr_OperatingExpensesAbstract	
• rr_ManagementFeesOverAssets	1
• rr_Component1OtherExpensesOverAssets	2
• rr_ExpensesOverAssets	3

Since the presentation linkbase does not include the rr_ProspectusShareClass element as a source so as to provide an order for its targets, the columns might be displayed in random order, this way:

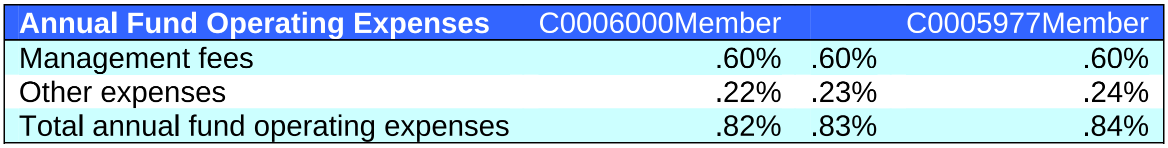

Annual Fund Operating Expenses	C0006000Member		C0005977Member
Management fees	.60%	.60%	.60%
Other expenses	.22%	.23%	.24%
Total annual fund operating expenses	.82%	.83%	.84%

Now, assume this presentation group:

Element	**Order**
rr_OperatingExpensesAbstract	
• rr_ManagementFeesOverAssets	1
• rr_OtherExpensesOverAssets	2
• rr_ExpensesOverAssets	3
rr_ProspectusShareClassAxis	
• rr_ShareClassDomain	1
• eg_S0006000	2
• eg_C0005977	3

The presentation group not only orders the children of the abstract element; it also orders the three contexts according to the children of rr_ProspectusShareClassAxis, resulting in a table like this:

Annual Fund Operating Expenses		C0006000Member	C0005977Member
Management fees	.60%	.60%	.60%
Other expenses	.22%	.23%	.24%
Total annual fund operating expenses	.82%	.83%	.84%

Only the immediate children of the axis element matter for the ordering of the contexts (and therefore the facts in the context). There is no indentation of the sort that can be done with the primary concepts' presentation hierarchy.

The presentation group has no effect on ordering with respect to the unit or language of a fact, and has no effect with respect to the identifier or period of a context.

Although the default member of an axis cannot appear in the xbrli:segment element of any context, it is important that the default member appear in the presentation group: this allows the preparer to control whether facts having the default member of an axis to appear at the beginning, end, or anywhere in a sequence of domain members. For an example, see example "eggplant" and role http://xbrl.sec.gov/rr/role/PerformanceTableData in which element rr_PerformanceMeasureAxis has six children, the fourth of which is the default member, rr_PerformanceMeasureDomain.

5.3 *Preferred Labels for Row and Column Headings*

The headings shown for a row or column where a primary concept appears depend on the preferred label attribute of the relationship between the concept's parent and the primary concept itself. The preferred label, if not present, defaults to the "standard" element label.

The preferred label attribute value of the relationship between the axis element and the domain member now also determines the heading shown for a row or column where a domain member of an axis appears.

The default member of an axis is a special case; it is treated as if its standard label were empty. If another label is desired for the default member, then use "terseLabel", "verboseLabel", etc.

For example, assume this presentation group:

Element	Preferred label Attribute value	Preferred label Text value	Order
rr_OperatingExpensesAbstract			
• rr_ManagementFeesOverAssets	label	Management fees	1
• rr_OtherExpensesOverAssets	terseLabel	Other expenses	2
• rr_ExpensesOverAssets	totalLabel	Total Expenses	3
rr_ProspectusShareClassAxis	label	Prospectus Share Class [Axis]	
• rr_ShareClassDomain	terseLabel	***	1
• eg_S0006000	verboseLabel	Share Class A	2
• eg_S0005977	terseLabel	B	3

Given the same data as above, the preferred labels have this effect:

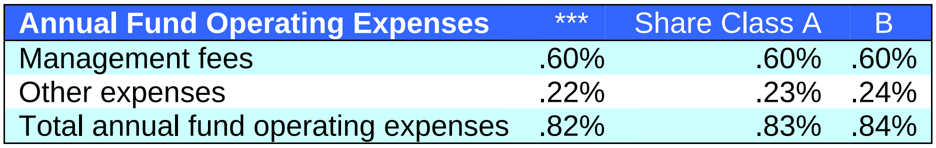

Annual Fund Operating Expenses	***	Share Class A	B
Management fees	.60%	.60%	.60%
Other expenses	.22%	.23%	.24%
Total annual fund operating expenses	.82%	.83%	.84%

If an axis is the source of a relationship with a preferred label that is a negating label role, there is no additional effect on any numeric facts displayed in that column or row. Put another way, negating of numeric facts depends entirely on the preferred label of a *primary* concept.

5.4 Details of context sorting and filtering

Only effective presentation relationships are relevant.

The presence of “hypercube” elements in parent-child relationships is not relevant.

Only parent-child relationships with an axis as source and a domain member as target have an impact on filtering contexts.

If axes are ordered relative to each other with some common parent, that has no impact on their sorting or filtering.

The Interactive Data being rendered is assumed to be valid with respect to EDGAR Filer Manual Volume II chapter 6:

- If more than one relationship in a presentation group have the same source and same value for the order attribute, the behavior is undefined.
- If there is more than one fact having the same primary element and equivalent context and units, then the behavior is undefined.
- If there are contexts with different values of element xbrli:identifier, then the behavior is undefined.
- If the xbrli:scenario element is present in a context, then the behavior is undefined.
- If a fact appears in a tuple, then the behavior is undefined.

- If a fact is not displayed in any presentation groups, then it is displayed in an “uncategorized items” report.

6 Embedding Commands in Text Blocks

Conventional rendering of Interactive Data consists only of one report per presentation group, ordered by the dictionary sorting of the presentation group definition text, and containing all the facts whose primary item appears in that presentation group. These are called the “top level” reports. Embedding commands allow additional reports that are not at top level.

A text block fact may contain an embedding command, which is a section of XHTML that looks like the following, anywhere in the text block:

```
&lt;div style="display: none">
~ http://xbrl.sec.gov/rr/role/ShareholderFeesData
column period compact *
row dei_LegalEntityAxis compact eg_S000005977Member
row dei_ProspectusShareClassAxis compact *
~ Comments can go here &lt;/div>
```

Because the un-escaped text is surrounded by the XHTML <div> element with style attribute “display: none”, even if the fact were to be rendered by an application that could not interpret the embedding command, the command would not be shown.

The purpose of the embedding command is to give the preparer the ability to place the usual report layout, restricted to certain line items, axis members, units, and periods, and to position it in place of where a text block could have gone instead.

It is neither necessary nor desirable for SEC filers creating interactive data instances to include additional XHTML markup that is not hidden, in the same text block. But, a preparer could do so, and the additional XHTML markup would then be displayed by applications that could not interpret the embedding command.

This feature in combination with context sorting and filtering allows rendering to produce layouts like that shown below:

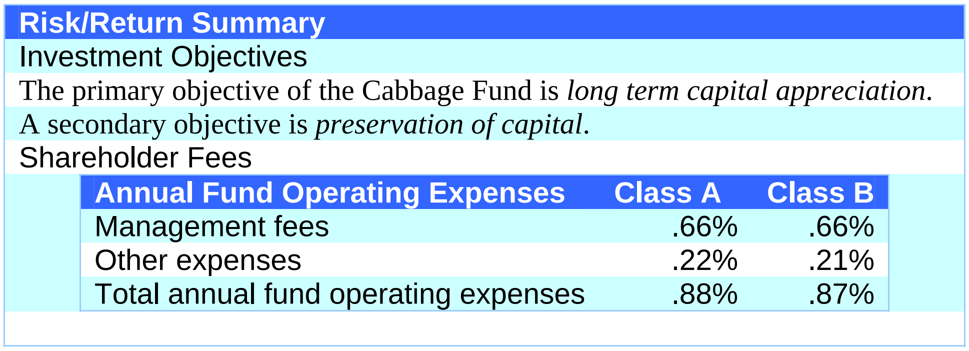

Risk/Return Summary

Investment Objectives

The primary objective of the Cabbage Fund is *long term capital appreciation*.

A secondary objective is *preservation of capital*.

Shareholder Fees

Annual Fund Operating Expenses	Class A	Class B
Management fees	.66%	.66%
Other expenses	.22%	.21%
Total annual fund operating expenses	.88%	.87%

6.1 Syntax of an Embedding Command

An embedding command consists of a role followed by a series of *iterators*. (Readers familiar with XBRL Formula Linkbases will recognize that an iterator resembles a filter).

An iterator consists of tokens separated by whitespace; the tilde character marks the beginning and end of the command.

- The beginning and ending tilde (~) indicate that the content of the enclosing element is an embedding command. Any text appearing before the first or after the second tilde is ignored.
- The first token is the URI of a Presentation Group used to display the facts.
    - For example, “http://xbrl.sec.gov/rr/role/ShareholderFeesData”.
- The token “column” starts an iterator for what is to appear in a column heading.
- The token “row” starts an iterator for what is to appear in a row heading.
- Any token with an underscore (“_”) denotes an element:
    - The string after the underscore must be nonempty and is an element name.
    - The string before the underscore must be nonempty and it is a preferred namespace prefix as defined in the EDGAR Filer Manual volume II subsection 6.7.7.
        - For example, token “dei_LegalEntityAxis” denotes the element “LegalEntityAxis” in the namespace with preferred prefix “dei” (usually “http://xbrl.us/dei/2009-01-31”).
- Within each iterator,
    - The first token may be “period”, “primary”, or a token with an underscore.
    - The second token is the word “compact”.
    - The third token is either “*” meaning “all”, or a token with an underscore that denotes a domain member.
- Column iterators must appear before row iterators.
- The order of the iterators determines their nesting, and is therefore significant.
- If the command has no iterator for “period”, it is assumed as an iterator starting the entire command, “column period compact *”.
- If there are facts using an axis for which there is no iterator, the iterator implicitly appears before any other columns as “column {axis} compact *”. The nesting of multiple implicit axis iterators is unpredictable.
- There can be any number of column iterators.
- If the command has no iterator for “primary”, it is assumed that the entire command ends with iterator “row primary compact *”.
- There can be any number of row iterators.

For example, suppose an instance has facts in contexts using the Legal Entity axis. This embedding command:

```
~ http://eggplant.com/role ~
```

would be interpreted as if it had been written:

```
~ http://eggplant.com/role
column period compact *
column dei_LegalEntityAxis compact *
row primary compact * ~
```

Conventionally, layout and rendering information has been placed entirely in the presentation linkbase. However, in a typical Interactive Data filing, the instance already contains a large amount of rendering information, but the information is encoded rigidly as XHTML in text blocks. In the embedding approach, those text blocks are made more flexible and can contain just enough information to connect a variety of facts and contexts with its presentation linkbase. The presentation linkbase still contains the general and reusable layout information.

6.2 *Columns, Rows and Nesting with Embedding Commands*

The embedding command is processed with column iterator 1 being the outermost grouping of columns, column iterator 2 nested inside those, and column iterator 3 inside iterator 2 and so on.

Similarly, row iterator 1 is the outermost grouping, followed by row iterator 2, and so on.

Suppose for example that we have a presentation linkbase with one abstract and one non-abstract row of "primary" Shareholder Fees data. There is a LegalEntityAxis with two members: the default and eg_S0005977Member. There is a ProspectusShareClassAxis and it has three members: a default, eg_C000001 and eg_C00002. At most there are six distinct contexts:

dei:LegalEntityAxis – empty (default) and rr_ProspectusShareClassAxis empty (default)
dei:LegalEntityAxis – empty (default) and rr_ProspectusShareClassAxis C00000001
dei:LegalEntityAxis – empty (default) and rr_ProspectusShareClassAxis C00000002
dei:LegalEntityAxis – S0005977 and rr_ProspectusShareClassAxis empty (default)
dei:LegalEntityAxis – S0005977 and rr_ProspectusShareClassAxis C0000001
dei:LegalEntityAxis – S0005977 and rr_ProspectusShareClassAxis C0000002

For the sake of illustration only, assume there is a "Distribution and/or service (12b-1) fees" percentage reported for each of these six combinations including defaults.

Given this command:

```
~ http://xbrl.sec.gov/rr/role/OperatingExpensesData
column period compact *
column rr_ProspectusShareClassAxis compact *
row dei_LegalEntityAxis compact *
row primary compact * ~
```

The resulting table is displayed this way:

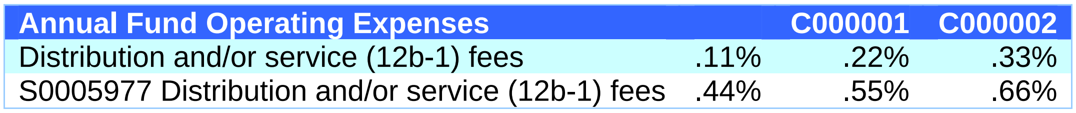

Annual Fund Operating Expenses		C000001	C000002
Distribution and/or service (12b-1) fees	.11%	.22%	.33%
S0005977 Distribution and/or service (12b-1) fees	.44%	.55%	.66%

Given this command instead:

```
~ http://xbrl.sec.gov/rr/role/OperatingExpensesData
column period compact *
column dei_LegalEntityAxis compact *
column rr_ProspectusShareClassAxis compact *
row primary compact * ~
```

The resulting table is displayed with this nesting, with legal entity outside of share class:

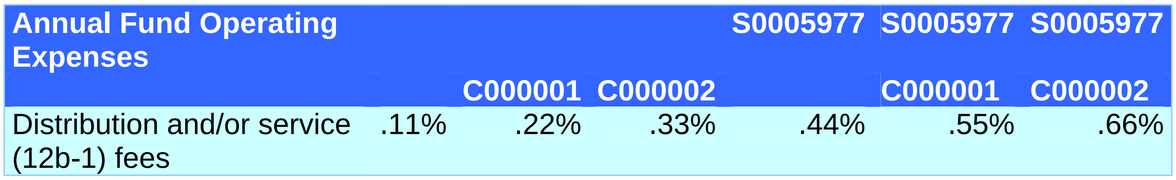

Annual Fund Operating Expenses				S0005977	S0005977	S0005977
		C000001	C000002		C000001	C000002
Distribution and/or service (12b-1) fees	.11%	.22%	.33%	.44%	.55%	.66%

For more elaborate examples of nesting, see the example files.

6.3 *Filtering with Embedding Commands*

If the token that appears in an axis iterator after the word "compact" is other than the token "*", then the token denotes a domain member. Only facts in contexts with that domain member in the specified axis will appear. For example, given this command with the same data:

```
~ http://xbrl.sec.gov/rr/role/OperatingExpensesData
column period compact *
column dei_LegalEntityAxis compact *
column rr_ProspectusShareClassAxis compact eg_C000002Member
row primary compact * ~
```

Then the display will appear as

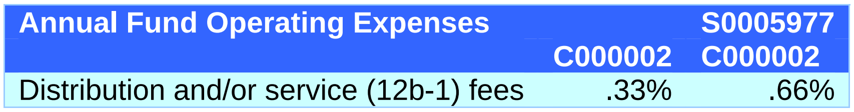

Annual Fund Operating Expenses		S0005977
	C000002	C000002
Distribution and/or service (12b-1) fees	.33%	.66%

Use the name of the default member to specify that only facts with the default in that axis will appear. For example, this command specifies the default member of one axis:

```
~ http://xbrl.sec.gov/rr/role/OperatingExpensesData
column period compact *
column dei_LegalEntityAxis compact *
column rr_ProspectusShareClassAxis compact rr_ProspectusShareClassDomain
row primary compact * ~
```

This embedding command produces the result

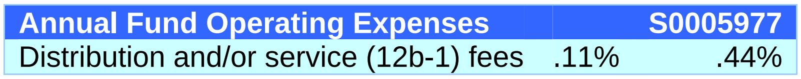

Annual Fund Operating Expenses		S0005977
Distribution and/or service (12b-1) fees	.11%	.44%

The filtering that happens from the iterator logically occurs before the filtering imposed by the presentation linkbase (presentation linkbase filtering occurs whether the Report is generated at top level or via embedded command). Being able to filter in either or both ways provides the preparer with fine grained control over what appears in individual tables, with one presentation linkbase potentially referred to in several different embedding commands.

For more elaborate examples, see particularly "cabbage", in which almost all of the embedding commands have an iterator restricting the data to a specific fund series.

6.4 Embedding and Flow-Through suppression

The term “flow-through suppression” refers to the feature of the SEC Viewer in which a column of facts that has already been presented in a report will not be presented again.

Flow-through suppression also applies to reports created from embedded commands.

For example, suppose there are two presentation groups 001 and 002. An instance has a text block in presentation group 001, and it contains an embedded command that displays all the facts in presentation group 002. Therefore, only presentation group 001 will be displayed at top level.

Flow-through suppression does not apply to a report with the {Elements} token. An {Elements} report displays all its facts.

If a preparer does nothing, flow-through suppression will usually ensure that each fact is displayed in one location. Sometimes it is important for the same facts to be displayed in multiple tables. Preparers can do this simply by using more than one embedded command for the same presentation group and filters.

The example “eggplant” shows how a single series’ data could be shown in multiple prospectuses within a single instance by using dei:DocumentInformationDocumentAxis. This is not an exception to the “flow through suppression”, because the one series’ data is displayed three times in different embedded tables.

6.5 Embedding and Footnotes

Footnotes of a report are displayed at the bottom of a report, sequentially numbered, and the footnote marks being numbers that appear to the right of each fact to which they refer. The SEC Viewer implements a new enhancement:

- If all of the facts on a row have the same footnote mark, then that footnote mark is promoted to a position just to the right of the row labels.
- If all of the facts on a column have the same footnote mark, then that footnote mark is promoted to the heading of the column.

For example, instead of this:

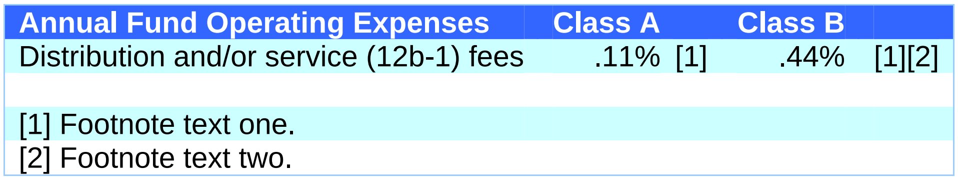

Annual Fund Operating Expenses	Class A		Class B	
Distribution and/or service (12b-1) fees	.11%	[1]	.44%	[1][2]
[1] Footnote text one.				
[2] Footnote text two.				

This would display as:

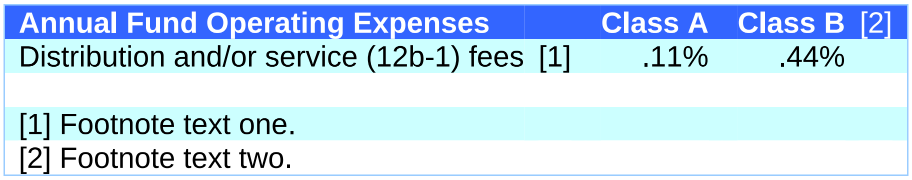

Annual Fund Operating Expenses		Class A	Class B [2]
Distribution and/or service (12b-1) fees	[1]	.11%	.44%
[1] Footnote text one.			
[2] Footnote text two.			

XBRL Footnotes are displayed at the bottom of a report even if the report is embedded.

The conventional notion of Interactive Data rendering is based on a complete financial report. Some SEC filings, such as filings pursuant to Rule 497 under the Securities Act of 1933, consist of material that may change a previous filing. This means that sometimes a filing may contain a table with footnotes numbered 1 and 2, and another table with footnotes numbered 3 and 4; but a subsequent filing that amends only the second of these tables would have footnotes numbered starting at 3. The SEC Viewer does not support this; every report's footnote numbering is local to that report. Example "dandelion" shows filers how to handle this. It includes the minimum text from the original, but the display of footnotes differs from the original.

The Interactive Data being rendered is assumed to be valid with respect to EDGAR Filer Manual Volume II chapter 6:

- If the footnote uses a custom footnote arcrole or role, then the behavior is undefined.

7 Risk/Return Summary Special Cases

7.1 Embedded Bar Chart Graphic

SEC Form N-1A requires the display of bar chart data graphically.

When an embedded command has a role URI that contains the phrase "BarChart" and the primary elements are drawn from the set rr_AnnualReturn1990, rr_AnnualReturn1991, and so forth, then the SEC Viewer displays a graphic instead of a conventional table.

7.2 Number Formats

The value 0 may be significantly different from a blank but may also need to be visually distinct from the value 0. For example, in the Shareholder Fees section of Form N-1A, the appearance of a 0 is conventionally indicated as the word "none".

There is no syntax for indicating this at the level of a role declaration or embedding command. Instead, specific Mutual Fund Risk/Return Summary Taxonomy 2010 elements are treated specially, as follows:

Element	ElementTypeName	How to display Zero	How to display Nonzero
rr_ExchangeFee	NonNegativeMonetaryType	None	
rr_MaximumAccountFee	NonNegativeMonetaryType	None	
rr_ShareholderFeeOther	NonNegativeMonetaryType	None	
rr_ExpenseExampleYear01	NonNegativeMonetaryType	None	
rr_ExpenseExampleYear03	NonNegativeMonetaryType	None	
rr_ExpenseExampleYear05	NonNegativeMonetaryType	None	
rr_ExpenseExampleYear10	NonNegativeMonetaryType	None	
rr_ExpenseExampleNoRedemptionYear01	NonNegativeMonetaryType	None	
rr_ExpenseExampleNoRedemptionYear03	NonNegativeMonetaryType	None	
rr_ExpenseExampleNoRedemptionYear05	NonNegativeMonetaryType	None	
rr_ExpenseExampleNoRedemptionYear10	NonNegativeMonetaryType	None	
rr_MaximumCumulativeSalesChargeOverOfferingPrice	NonNegativePure4Type	None	%
rr_MaximumCumulativeSalesChargeOverOther	NonNegativePure4Type	None	%
rr_MaximumSalesChargeImposedOnPurchasesOverOfferingPrice	NonNegativePure4Type	None	%
rr_MaximumDeferredSalesChargeOverOther	NonNegativePure4Type	None	%
rr_MaximumDeferredSalesChargeOverOfferingPrice	NonNegativePure4Type	None	%
rr_MaximumSalesChargeOnReinvestedDividendsAndDistributionsOverOther	NonNegativePure4Type	None	%
rr_ExchangeFeeOverRedemption	NonNegativePure4Type	None	%
rr_MaximumAccountFeeOverAssets	NonNegativePure4Type	None	%

Element	ElementTypeName	How to display Zero	How to display Nonzero
rr_ManagementFeesOverAssets	NonNegativePure4Type	None	%
rr_DistributionAndService12b1FeesOverAssets	NonNegativePure4Type	None	%
rr_DistributionOrSimilarNon12b1FeesOverAssets	NonNegativePure4Type	None	%
rr_OtherExpensesOverAssets	NonNegativePure4Type	None	%
rr_Component1OtherExpensesOverAssets	NonNegativePure4Type	None	%
rr_Component2OtherExpensesOverAssets	NonNegativePure4Type	None	%
rr_Component3OtherExpensesOverAssets	NonNegativePure4Type	None	%
rr_AcquiredFundFeesAndExpensesOverAssets	NonNegativePure4Type	None	%
rr_ExpensesOverAssets	NonNegativePure4Type	None	%
rr_NetExpensesOverAssets	NonNegativePure4Type	None	%
rr_RedemptionFeeOverRedemption	NonPositivePure4Type	None	%
rr_FeeWaiverOrReimbursementOverAssets	NonPositivePure4Type	None	%
ALL OTHERS	*ALL OTHERS*	0	

8 Technical End Notes

8.1 Recommended Namespace Prefixes

In an XBRL instance the namespace *prefix* is arbitrary; it is the *namespace* of the elements that matters. For example, this element would have the exact same semantics as one with conventional namespace prefixes:

```
<n1:segment
xmlns:n1="http://www.xbrl.org/2003/instance"
xmlns:n2="http://xbrl.org/2006/xbrldi"
xmlns:n3="http://xbrl.us/dei/2009-01-31"
xmlns:n4="http://cabbage/20100501"
><n2:explicitMember
      dimension="n3:DocumentInformationDocumentAxis"
      >n4:RetirementMember</n2:explicitMember></n1:segment>
```