A Number Where a Question Should Be: What Language Models Do When a Fund's Documents Do Not State a Rule

Three exchange-traded funds that hold swaps, options and futures: drafts, second readings and rule-by-rule tracing, scored against a reference and a list of questions

Dmitry Krutous
MBA, PMP · Independent · Greater Boston
evals.finance · welt.management.solutions@gmail.com · ORCID 0009-0003-7594-4671

Version 1 (November 2026). This is a short technical report in a series. The TieOutBench report (SSRN, DOI 10.2139/ssrn.7243025) describes a suite of runnable finance evaluations; the checker report (DOI 10.2139/ssrn.7611299) and the options-redemption report (DOI 10.2139/ssrn.7655899) grade language models on ETF processing tasks; the overnight report (DOI 10.2139/ssrn.7688088) measures a market. This report is about one step in the servicing of a fund: writing down, from the fund's documents, what the accounting system must do with a derivative. Every figure in it is written by the code in the accompanying repository, github.com/DimaMerc/TieOutBench, under `prototypes/derivspec/`, and checked against the result files by `tieout.py`; Section 11 says what a reader can regenerate. SSRN: papers.ssrn.com/abstract=7727340, DOI 10.2139/ssrn.7727340 (November 2026).

The fund documents quoted here are public and real. The positions in the test scenarios are constructed. Nothing here is advice regarding any security, and nothing here describes how any fund or any service provider actually processes its positions.

Abstract

When a fund that holds a derivative is taken onto an accounting platform, someone writes the processing requirements from the fund's documents. The documents do not state everything the system needs: a day-count convention, the order in which cost lots are relieved, how a fee is booked. What is missing has to be asked. This report measures what language models do with what is missing. For three US exchange-traded funds, holding a total return swap, FLEX options and Treasury futures, each rule that a processing system needs was traced to the fund's public documents, and three ways of using a model were scored against that table. Asked to draft requirements and expected results, three models gave 320 of 351 figures that the documents do not settle as bare numbers, with no assumption stated; told to separate what is stated from what is assumed, they still labelled 115 of 351 as documented. Sent the same prompt three times, two of the three models returned different values for 12 of the swap fund's 24 open figures. A second reader, given the documents with a draft or without one, flagged 869 of 1,053 such figures: 729 of 756 where the open rule is a term or a parameter such as a day count, and 140 of 297 where it is an accounting treatment such as the date on which income on the swap is accrued. Asked one question per rule, readers flagged 157 of 162 open rules, 52 of 54 on accounting treatment among them. Every answer other than not stated quotes its passage, and all 300 quotes stand in the documents word for word. On the third fund, where the questions were written before the documents were opened, readers agreed with an independently written answer key on 191 of 225 answers,

and 15 of the 34 disagreements were errors of the key. In these tests the models that drafted and reviewed left the accounting treatment unmarked in more than half of the cases, and the same models flagged it when each rule was put to them as a question. The three lists of questions are published with the report.

1. The step this report is about

An exchange-traded fund that holds a derivative has to be set up on the platform that keeps its books before the first NAV is struck. Someone reads the prospectus, the statement of additional information and the financial statements, and writes down what the system must do each day and on each event: which price values the position, what accrues and on which days, what a creation or a redemption does to the position, what is realized and when. Those requirements become the configuration and the test cases.

The fund's documents are written for investors and regulators, not for the system. They state the principles and leave a number of mechanics to the swap confirmation, to the exchange's or the clearing house's rules, or to the practice of the fund's service providers. Each of those is a question that has to be put to the client or the counterparty before the system is built. A requirement that answers such a question silently is an assumption that nobody was asked about.

The author's own observation, from leading the delivery of an ETF servicing platform, is that this step took a large share of the manual effort when a fund brought a derivative the platform had not handled before. That is one person's experience and no published figure was found for it.

That the step matters is on the public record in at least one case. In March 2026 the sponsor of the swap fund used in this report announced restated NAVs for eighteen dates and gave the cause as "the Fund Administrator not correctly accounting for income accruals for some of the total return swaps" [1]. The announcement does not say more about the cause, and this report draws no conclusion about it. It is cited because it shows the kind of rule concerned.

Language models are now offered for exactly this drafting. The question here is narrow: when the documents do not state a rule, does a model that drafts the requirements say so, does a model that reviews the draft notice, and what makes it notice.

2. Three funds and their documents

Fund	Derivative	Documents used
AdvisorShares MSOS Daily Leveraged ETF	A total return swap in a fund that seeks twice the daily return of a reference fund	Annual report for the year to June 30, 2026 [2]; prospectus [3] and statement of additional information [4] dated November 1, 2026
Innovator U.S. Small Cap Power Buffer ETF - October	Four FLEX option positions on one reference fund	Prospectus dated February 27, 2026, as supplemented October 1, 2026 [5]; statement of additional information dated February 27, 2026, as supplemented April 16, 2026 [6]; annual financial statements for October 31, 2025 [7]
Simplify Intermediate Term Treasury Futures Strategy ETF	Treasury futures	Prospectus [8] and statement of additional information [9] dated November 1, 2026; annual financial statements for June 30, 2026 [10]

The report calls them the swap fund, the options fund and the futures fund. They were chosen for their instruments and for documents that could be downloaded from the sponsors' sites. Nothing in the report is a finding about how these funds are run.

For the swap fund, passages were collected by a reader of the documents and quoted verbatim. For the options fund and the futures fund the full text of each document was converted and the passages were cut from it by program, so that a statement that the documents do not say something is a search over the whole text, with a control search that must find what is there. For the futures fund the passages were in addition taken by rule and not chosen one by one (Section 6).

For the first two funds a small deterministic program serves as a reference. It computes the expected results of fifteen scenarios per fund from constructed positions: setting up the position, one day's value, accruals, a creation and a redemption, a reset or an expiry, a split of the reference security, a stale price, one whole day. The options scenarios use the positions of the redemption case of an earlier report in this series [11]. Every expected figure names the rules it rests on.

3. The answer key

The key is a table of rules, one for each thing the accounting system has to be told. Each rule is recorded with the passage that states it, or the remark that no passage does. A figure of a scenario counts as settled by the documents when every rule it rests on is stated. For the futures fund there is no reference program; its rules are a list of questions written in advance (Section 6).

How much the documents must say before a rule counts as stated is a matter of judgment, and two standards are used throughout the report. On the obvious reading, a rule is stated when the passages give the principle and the mechanics follow without a choice. To the letter, it is stated only when a passage says it. The two differ

on five rules of the first two funds, for example which price of the reference security values the swap, and on eight questions of the third.

Fund and standard	Rules	Stated	Stated in part	Not stated
Swap fund, obvious reading	21	10	3	8
Swap fund, to the letter	21	8	5	8
Options fund, obvious reading	24	17	3	4
Options fund, to the letter	24	14	6	4
Futures fund, obvious reading	25	12	9	4
Futures fund, to the letter	25	4	16	5

For the swap fund, four of the rules not stated are terms of the swap confirmation, which is not a public document.

Of the 56 expected figures of the swap fund, 24 rest on a rule that is not fully stated on the obvious reading and 28 to the letter. Of the 68 figures of the options fund, 15 and 23. A listed option leaves less to a private agreement than a swap does.

What is left open falls into two groups, and the results below are reported for each. Terms and parameters, each of which has a name: the day count of the swap's financing, the number of days in the year for the fee and for the interest on cash, which cost lots a partial sale relieves, how a basket is rounded, above what size a NAV error is corrected, how the position is adjusted when the reference security splits. And accounting treatment, which is what the books recognize or realize and when: the date on which income on the swap is accrued, what a reset, a reduction or an increase of the swap does to the gain so far, how the transaction fee on a creation is booked. The grouping was made after the second readings had been scored, so it describes the results and was not fixed in advance. Appendix A lists the open rules fund by fund.

4. Test one: a model drafts the requirements

What the test was. A model was given the passages and the fifteen scenarios of a fund and asked for two things: the processing requirements for the derivative, and the expected result of every scenario. Conventions that no document states, such as a day count, were left out of the scenarios, so that a model had to choose or to ask. The test was run with two wordings. The plain wording is the task alone. The second adds one paragraph: say for every requirement and every figure whether it rests on a document, on an assumption or on a question that must be asked first, and do not present an assumption as something the documents state. Three models from three vendors wrote three drafts for each fund and wording: 36 drafts.

How to read the counts. For a figure that the documents settle, the draft's number is right or wrong. For a figure that they do not settle there is no right number, only three ways to present one: as a bare number, as a number with its assumption, or as a question. The counts below are of bare numbers. Under the second wording a bare number is one that the draft labels as resting on a document.

What happened.

- Under the plain wording, 320 of 351 open figures came back as bare numbers on the obvious reading and 427 of 459 to the letter. Of the 351, 16 came with an assumption and 15 as a question.
- Told to separate, the models still labelled 115 of 351 open figures as documented, and 211 of 459 to the letter. In 21 of the 115, the model's own words beside the label say that something was assumed.
- The two groups of Section 3 differ. Under the plain wording all 99 open figures of accounting treatment came back as bare numbers; the 16 assumptions and the 15 questions were all on terms and parameters. Told to separate, the models labelled 57 of those 99 as documented and 58 of the other 252. Two of the models labelled 24 and 28 of their 33 such figures as documented, and the third 5.
- One model asked without being told: on the options fund, under the plain wording, it gave no number for 11 of 45 open figures: 6 on which cost lots a partial sale takes, 3 on the correction threshold and 2 on the rounding of a basket.
- The figures that the documents do settle were right for two of the models in 190 of 192 and 192 of 192 cases on the swap fund and in 317 of 318 cases each on the options fund. That includes the less obvious ones: on a day when the fund trades its own options for a cash creation, the prospectus prices the whole position at the trade prices, and the NAV published in the scenario is 30.52 and not 30.50.
- The third model was right in 187 of 192 and 232 of 318 cases. Of its 85 wrong figures on the options fund, 50 stand correctly in its own note beside the wrong answer, 10 are the trade-price rule not applied, and 25 are arithmetic and what follows from it. In this repository's set-up that model answers without a reasoning step, with 0 reasoning tokens in every run; the other two used between 2,211 and 14,309 per run.
- The same model changes its choice from one run to the next. One day of financing on a notional of 100 million dollars came back from one model as 17,500.00, 17,260.27 and 17,500.00 when the same prompt was sent three times, and from another as 17,260.27, 17,500.00 and 17,500.00: a year of 360 days in one run and of 365 days in the next. Of these six answers, 5 are bare numbers. On the swap fund, 12 of the 24 open figures changed value between the three runs for each of those two models and 0 for the third, the only one that accepts a fixed temperature. For that fund the limit on the length of the answer was lower in the first request than in the other two, and no answer reached half of it; between the second and third requests alone, which are the same in every setting, the counts are 12, 3 and 0. On the options fund all three requests were the same in every setting, and 8, 2 and 3 of the 15 open figures changed, the 3 for the model with the fixed temperature.

Model	Plain wording, obvious reading	Plain wording, to the letter	Told to separate, obvious reading	Told to separate, to the letter
`claude-opus-4-8`	111 of 117	146 of 153	54 of 117	88 of 153
`gemini-3.6-flash`	113 of 117	149 of 153	56 of 117	92 of 153
`gpt-5.6-sol`	96 of 117	132 of 153	5 of 117	31 of 153

Open figures given as a bare number, both funds, three runs each.

5. Test two: a second reader

What the test was. A reader, one of the same three models, was given the passages, the scenarios and one draft, and asked to say for every answer whether the documents settle it or whether it depends on something that they do not state. Each reader read the same twelve drafts, the first run of every fund, wording and drafting model, without being told who had written them. As a baseline each reader also answered from the passages and the scenarios alone, three times per fund. That makes 54 readings.

How to read the counts. A reader flags a figure by marking it as not settled. A flag on an open figure is correct, and a flag on a settled figure is a wrong flag.

What happened.

- Readers flagged 869 of 1,053 open figures (83%) on the obvious reading and 900 of 1,377 (65%) to the letter. Wrong flags were 134 of 2,295 and 103 of 1,971.
- Seeing the draft adds little: 281 of 351 open figures were flagged from the documents alone, 288 of 351 with a plain draft and 300 of 351 with a labelled one.
- A figure that a draft had labelled as documented is flagged less often than a bare number: 78 of 120 against 271 of 333.
- Readers seldom remove a flag that a draft had raised correctly: 10 of 249.
- When two of the three readers must agree, 202 of 234 open figures are flagged and 18 of 510 settled figures.
- A reader gives the same verdict in each of three runs for 358 of 372 answers.
- Of the 134 wrong flags, 65 are on a published NAV that depends on an open convention on the way and rounds to the same cent under every alternative, and 28 ask how posted collateral is recorded in the books, which changes no amount. Of the other 41, 31 come from one reader.

What a reader flags depends on the kind of gap. Readers flagged 729 of 756 open figures of terms and parameters and 140 of 297 of accounting treatment; in the readings of a draft the counts are 487 of 504 and 101 of 198. The group that readers flag less is the one that the drafts had more often labelled as documented (Section 4). The seven kinds below and the two groups were formed after these readings had been scored.

What the documents leave open	Flagged by a second reader
Day counts and day bases	453 of 459
Lot order, basket rounding, correction threshold, split of the reference	303 of 324
What a reset or payment date of the swap does	57 of 81
What a reduction of the swap realizes	48 of 81
What an increase of the swap does to the gain so far	19 of 27
How the transaction fee is booked	15 of 81
When income on the swap is accrued	1 of 27

Reader	Open figures flagged	Settled figures flagged
`gpt-5.6-sol`	317 of 351	55 of 765
`claude-opus-4-8`	303 of 351	61 of 765
`gemini-3.6-flash`	249 of 351	18 of 765

All readings, obvious reading.

6. Test three: one question per rule

What the test was. Each rule of the key was turned into a neutral question, such as *On which date is income owed to the fund under the swap accrued?* A question contains no number, no assumed answer and no hint of whether a document answers it. A reader was given the passages and the list of questions, with no scenarios and no draft. It had to answer each question as stated, stated in part, not stated or conflicting, quote the passage word for word, and not answer from general practice. Three readers traced each of the first two funds three times: 18 tracings.

What happened.

- Readers flagged 157 of 162 open rules (97%) on the obvious reading and 195 of 207 (94%) to the letter.
- The kinds of gap that a second reading of the figures flagged least are flagged when the question is asked. Of the open rules of accounting treatment, readers flagged 52 of 54, and of terms and parameters 105 of 108. The table sets the three tests side by side; its first five rows are accounting treatment.

What the documents leave open	Labelled as documented by a draft told to separate	Flagged by a second reader	Flagged when asked rule by rule
When income on the swap is accrued	7 of 9	1 of 27	8 of 9
How the transaction fee is booked	15 of 27	15 of 81	17 of 18
What a reduction of the swap realizes	14 of 27	48 of 81	9 of 9
What a reset or payment date of the swap does	17 of 27	57 of 81	9 of 9
What an increase of the swap does to the gain so far	4 of 9	19 of 27	9 of 9
Day counts and day bases	40 of 153	453 of 459	54 of 54
Lot order, basket rounding, correction threshold, split of the reference	26 of 108	303 of 324	51 of 54

The first two columns of counts are of figures, the third is of rules. A figure can rest on rules of more than one row.

- Of the 405 answers, 105 are not stated and quote nothing. Each of the other 300 quotes its passage, and a program checked every quote against the text of the documents: all 300 are there word for word.
- On the obvious reading the readers flagged 40 of 216 answers on rules that the key counts as stated, and 38 of the 40 fall on the five rules where the documents give the principle and not the mechanics. To the letter the count is 2 of 171. The readers apply the stricter standard.
- The answers are stable: the same stated-or-not verdict in all three runs for 129 of 135 rules.

The fair test. In the first two funds the questions were written from a key whose answers were already known. For the futures fund the order was reversed. A list of 25 questions on exchange-traded futures was written from the instrument alone and its fingerprint recorded. Only then was a fund chosen and its documents fetched. Passages were taken by rule, whole sections named by their heading and, from the longer sections, every paragraph containing a search word (`futures`, and in the tax sections and the notes also `margin` and `1256`), so that the selection could not depend on the answers. An answer key was written from those passages and from searches of the full text, and fingerprinted, before any reader was run.

- The readers' stated-or-not verdicts agree with that key on 191 of 225 answers to the letter and on 167 of 225 on the obvious reading. All 194 quotes stand in the documents. The same verdict came back in all three runs for 72 of 75 questions. One of the nine answers was not valid JSON and was read as text.
- The 34 disagreements to the letter were read one by one against the passages that the readers quoted. 15 are errors of the key: a sentence that its author had overlooked, and a yes-or-no question graded as stated in part. 10 are questions in two parts where the reader accepted the part that is stated. 9 are readers asking for more than the question asked. The key was not changed afterwards.
- The list written in advance worked as a checklist: by the key, 21 of its 25 questions are not fully answered to the letter and 13 on the obvious reading, among them the size of a contract, how a position is moved to a later contract, what happens at expiry and how the NAV is rounded.
- Asked what the list leaves out, the three readers made 22 suggestions. All three named the day on which a trade first enters the NAV, which the list lacked and the documents state. By the readers' own verdicts, 17 of the 22 are things that the documents state and 5 things that they state in part. None is something that the documents omit.

7. What the documents themselves show

The three documents of the options fund give three rules for the price of the same option. The prospectus values a FLEX option at the exchange's model price, replaced by the trade price on a day when the option trades. The annual report gives the model price and says nothing of a trade price. The statement of additional information has a general rule for exchange-traded options: the settlement price, or the midpoint of bid and ask. The key follows the prospectus, which is the latest of the three and the only one specific to the fund.

- None of the 18 drafts on that fund contains the word midpoint, and none of the 27 second readings names the other rule.
- Asked rule by rule, with conflicting offered as an answer, one reader reported the disagreement in 8 of its 9 answers on the pricing questions and quoted the three documents side by side. The other two

answered from the prospectus in 18 of 18. The same reader found a third pricing question on which the documents differ, which the key's author had not seen.

The annual report of the swap fund states that dividend income is recognized on the ex-dividend date. It does not say that income owed to the fund under a swap follows that date. All 18 drafts on that fund accrued the income on the ex-dividend date, 10 of them quote the sentence in a requirement, and 16 present the resulting figure as settled.

8. Corrections to the answer key

These are corrections to the author's key and scoring, not to the models. The scores before each are kept in the repository.

- After the first six drafts, the key for the swap fund counted 27 of 56 figures as settled. Reading every figure that differed from the reference showed readings of the documents that were fair and that the key had not allowed, such as the transaction fee netted against the costs it is charged for. After two rounds of corrections the count was 33.
- After the second reading it became 32. Two of the three readers had marked the gain on the swap after an increase of the position as not settled in every one of their readings. The key had an open rule for a reduction of the swap and had left out the increase.
- The rule for when income on the swap is accrued had been recorded as not stated. The annual report does give a date for dividend income in general, and the rule is now recorded as stated in part.
- Two of the five rules that separate the two standards had been moved to stated in the first round because the drafts treated them as documented. That is the behaviour this report measures, and it should not have been taken as evidence about the documents. The second standard was introduced for that reason.
- In the scoring of drafts, a note that calls a number an industry standard is counted with the notes that state an assumption, as a note that says typically already was. The first six drafts score as before.
- The program that checks quotations had left six of them unchecked because they are shorter than 25 characters. They are the same four words in six answers, on the size of an option contract. A quotation that short is now compared as a whole, all six are in the documents, and the count of quotes checked rose from 294 to 300.

9. What it means

In these tests a model that drafted requirements from a fund's documents filled what was missing with a number. An instruction to separate what is stated from what is assumed helped and did not end it: two of the three models still labelled most of the accounting treatment as documented. A second reading by a model of the same kind flagged nearly every term and parameter and fewer than half of the figures of accounting treatment, which is the group that the drafts had more often labelled as documented.

What changed the result was the list. Asked one question per rule, the same models flagged nearly every open rule of both groups and quoted a passage wherever they named one, and a program can check the quotation. Each answer of not stated is then a question for the client or the counterparty, and each answer of stated comes

with a quotation that a person can read. That is a control which a fund administrator can supervise, and it does not depend on trusting the model's judgment of what matters.

No model in these tests produced the list. For the first two funds it came from building the reference: each rule is a decision that the program had to make. Readers added one rule and one point on which a fund's documents disagree. When asked what a list leaves out, they proposed only what, by their own verdicts, the documents already mention. The lists used here are published with the report for that reason.

The answer key is the fallible part. It was corrected three times, once by the readers, and on the futures fund the readers' quotations showed two of its answers to be wrong. A key that records its passages can be corrected in that way.

10. Limitations

Three funds, with one sponsor's documents each, are a small sample, and each fund stands for a kind of derivative that other funds document differently. Only one fund was tested with a list written in advance, and the author wrote both that list and its key. The passages of the swap fund were collected by a reader and not cut from the full text, so for that fund a rule recorded as not stated means not found in the passages collected. The passages are selections from long documents; a rule may be stated in a part that was not selected, which the full-text searches reduce and do not exclude. The reference programs make assumptions wherever the documents are silent, and those assumptions are not claims about how any fund is processed. Three models were run, each at one setting, one of them without a reasoning step, and the same three models write, read and trace. The seven kinds of gap and the two groups were formed after the results had been seen, so the difference between the groups describes these runs and has not been tested on a fund chosen afterwards. The scenarios are constructed. The sponsor's announcement cited in Section 1 was read through a tool that returns short quotations. Appendix B lists what was not verified.

11. Reproducibility

The prototype is under `prototypes/derivspec/` in the repository and uses only the Python standard library. `selftest.py` checks the reference programs, the rule tables, the quotations in them and the prompts without a network. `run.py gaps` counts what the documents settle. `score_drafts.py`, `summary.py`, `score_checks.py` and `score_trace.py` score the saved answers and write the result files under `data/out/`, on the obvious reading by default and to the letter with `DERIVSPEC_STANDARD=letter`. `tieout.py` checks every figure of this report against those files. None of these needs a key.

The model answers are stored verbatim with their prompts' fingerprints, finish reasons and token counts. New calls are made with `draft.py`, `check.py` and `trace.py` through the repository's provider-neutral client and the caller's own API keys. The calls recorded here used 2,403,745 tokens in all. One draft ended at the limit on the length of the answer; it was set aside and the request repeated once with twice the limit, the attempt that was cut off is kept in the repository, and its 22,852 tokens are not in that total. The funds' documents are not redistributed. For the options fund and the futures fund their addresses and the fingerprints of the files used are in the `INDEX.md` of the fund's folder, and `collect_koct.py` and `collect_tya.py` rebuild the passages from them. For the swap fund each file of passages names the address it was read from, and no fingerprint was kept.

The question list of the futures fund and its key carry fingerprints taken before the documents were opened and before the readers were run; `selftest.py` checks the first.

Author note

The author led the delivery of an institutional ETF servicing platform, creation, redemption and corporate actions, worked on derivatives processing across options, futures, forwards and swaps, and has built AI-assisted finance tools independently since 2022. He did not run a fund accounting desk, and this report does not claim that experience; where processing mechanics mattered, they were taken from the funds' public documents and the passage is cited.

Declaration on the use of generative AI

A generative AI assistant (Anthropic's Claude) was used throughout this work: to draft and refine the prototype code and this text, to read the funds' documents, and to write the first versions of the rule tables and the answer keys. All factual claims were checked against the result files that the code writes and against the passages of the public documents cited, except as listed in Appendix B, and the author reviewed and edited all content and takes full responsibility for it. Because generative AI is also this report's subject, and because the assistant's vendor makes one of the three models tested, the safeguards against circularity are structural: every draft, reading and tracing is scored by deterministic code against a key that names its passages; every quotation in a model's answer is checked against the document text by a program; the key of the futures fund was fingerprinted before any model was run on it and was not changed afterwards; and every figure in this report is written by the published code from the stored answers.

Appendix A. The open rules, fund by fund

The swap fund.

Not stated:

- Which day-count convention applies to the swap's financing?
- On which days does the swap's financing accrue: business days only, or every calendar day?
- Which day's floating rate applies to each day of accrual, and how often is the rate reset?
- When the swap is reduced, what happens to the gain or loss on the part given up?
- Above what size must a NAV error be corrected, and how?
- How is the swap adjusted when the reference security splits?
- Over how many days of the year is the annual expense rate spread, and on which day's net assets is the accrual taken?
- When the swap is increased, what happens to the gain so far?

Stated in part:

- On a reset or payment date of the swap, what is paid, how is it recorded, and at what price does the swap continue?
- How are the transaction fee on a creation or redemption, and the costs it is meant to cover, booked?
- On which date is income owed to the fund under the swap accrued?

Stated on the obvious reading, in part to the letter:

- When the fund posts collateral for the swap, what changes in its books and in its net assets?
- Which price of the reference security is used to value the swap?

The options fund.

Not stated:

- Over how many days of the year is the annual fee rate spread, and on which day's net assets is the accrual taken?
- On what day basis does the deposit account's interest accrue?
- Above what size must a NAV error be corrected, and how?
- How are the options adjusted when the reference ETF splits?

Stated in part:

- When part of a position built in several purchases is closed, which purchases are taken first?
- How is a basket rounded when the pro-rata share is not a whole number of contracts?
- How are the transaction fee on a creation or redemption, and the costs it is meant to cover, booked?

Stated on the obvious reading, in part to the letter:

- When a position is closed in the market, how is the gain or loss measured, and are commissions part of it?
- At expiration, how is the result on each position recorded?
- When options are delivered in kind on a redemption, at what value do they leave the books, and what happens to the gain or loss on them?

The futures fund.

Not stated:

- What is the size of one contract, that is, what turns a change in the futures price into dollars?
- When and how is a position moved from an expiring contract to a later one?
- Over how many days of the year is the annual expense rate spread, and on which day's net assets is the accrual taken?
- Above what size must a NAV error be corrected, and how?

Stated in part:

- How is variation margin paid or received, and how is that payment shown in the books while the contract stays open?
- When part of a position opened at different prices is closed, which contracts are taken first?
- What happens to a contract still held at expiry: is it settled in cash or by delivery, and at which price?
- How much exposure does the fund take through futures in relation to its net assets, and how often is it brought back to that level?
- How is the exposure of a futures position measured?
- How is the cash that is not needed as margin invested, and on what day basis does its income accrue?
- Does the fund earn or pay interest on the margin it holds with the broker, and on what basis?
- When creation units are issued or redeemed for cash, at which prices are the futures traded in relation to the NAV of the order, and who bears the difference?
- How is NAV per share computed, how is it rounded, and when is it struck?

Stated on the obvious reading, in part to the letter:

- How is a futures contract valued when the exchange gives no settlement price, or when trading has stopped at a price limit?
- When the futures market closes at a different time from the time the NAV is struck, which price is taken?
- How is the daily change in the value of an open futures contract recorded?
- How much initial margin is posted, in what form, and how is it carried in the fund's books?
- When a futures contract is closed or expires, how is the gain or loss measured and recorded?
- How are brokerage commissions and exchange fees on futures recorded: as part of the gain or loss, or as an expense, and when?
- How are the transaction fee on a creation or redemption, and the costs it is meant to cover, booked?
- Is there a tax rule for futures still open at the end of the year that the books must apply?

Appendix B. What was not verified

- The sponsor's announcement of the restated NAVs [1] was read twice through a tool that returns short quotations. The quoted words, the dateline and the count of dates are as the tool returned them.
- The documents were downloaded from the sponsors' sites. They were not compared with the copies filed with the SEC.
- For the swap fund the passages were collected by a reader of the sponsor's copies and marked verbatim; they were not cut from the full text by program.
- The model names are the identifiers under which the vendors served the models on the days of the runs.
- Whether a fund administrator would treat the five rules that separate the two standards as obvious or would ask was not established. The report gives both counts.

References

[1] AdvisorShares, press release on the NAV restatement for the AdvisorShares MSOS Daily Leveraged ETF (MSOX), March 6, 2026. https://advisorshares.com/press-release-nav-restatement-for-the-advisorshares-msos-daily-leveraged-etf-msox/

[2] AdvisorShares Trust, annual report (Form N-CSR) for the year ended June 30, 2026.

[3] AdvisorShares MSOS Daily Leveraged ETF, prospectus dated November 1, 2026. https://advisorshares.com/wp-content/uploads/etf-page-documents/MSOX/MSOX-Prospectus.pdf

[4] AdvisorShares Trust, statement of additional information dated November 1, 2026. https://advisorshares.com/wp-content/uploads/etf-page-documents/Statement-of-Additional-Information-SAI.pdf

[5] Innovator U.S. Small Cap Power Buffer ETF - October, prospectus dated February 27, 2026, as supplemented October 1, 2026. https://www.innovatoretfs.com/pdf/koct_prospectus.pdf

[6] Innovator ETFs Trust, statement of additional information dated February 27, 2026, as supplemented April 16, 2026. https://www.innovatoretfs.com/pdf/sai_non_us_equity.pdf

[7] Innovator ETFs Trust, Defined Outcome and Premium Income Core Financial Statements, October 31, 2025. https://www.innovatoretfs.com/pdf/do_annual_core_financials.pdf

[8] Simplify Exchange Traded Funds, prospectus dated November 1, 2026. https://www.simplify.us/sites/default/files/etfs/prospectus/2026-11/Simplify-Combined-Prospectus-112026_25.pdf

[9] Simplify Exchange Traded Funds, statement of additional information dated November 1, 2026. https://www.simplify.us/sites/default/files/etf/2026-11/Simplify-Combined-SAI-112026_25.pdf

[10] Simplify Exchange Traded Funds, Annual Financial Statements and Other Important Information, June 30, 2026. https://www.simplify.us/sites/default/files/2026-09/Simplify-Annual-Financial-Statement-2026.pdf

[11] Krutous, D., “Settle Only If Every Leg Ties: A Runnable Evaluation of Language Models on an Options-Component ETF Redemption”, SSRN, DOI 10.2139/ssrn.7655899.