



# How to Fix the FRM

**To:** FRM Committee members and GARP's Board of Trustees

**From:** David Harper, CFA, FRM, CIPM; Founder and CEO, bionicturtle.com

**Date:** 5/22/2018

**Subject:** Reputational Risks to the FRM Product; Related Questions; and Recommendations

---

## Contents

- Introduction
- GARP's Official Practice Exams exhibit a high error rate and send a poor signal (page 3)
- Evidence suggests actual exams have not been sufficiently robust and fair (page 7)
- The syllabus may not be receiving enough attention (page 8)
- Questions for the Board and the FRM Committee (page 13)
- Some recommendations and suggestions (page 14)
- Conclusion (page 15)
- *Appendix A: **Specific Questions** (“Short List”) for the Board and Trustees* (page 16)
- *Appendix B: My “Response to Bill May’s Response” [to the shared *preview* of this memo]* (page 17)

## Introduction

Hello, my name is David Harper. My company, Bionic Turtle LLC, has been an approved FRM Exam Preparation Provider (EPP) for over a decade. We have made several contributions to the FRM product<sup>1</sup>. Most of these contributions originated in customer interactions with FRM candidates, often when they are stuck in some way. For example, years ago we noticed that some candidates were having an excessively difficult time with certain problems merely because the compound frequency assumptions vary by author (e.g., continuous in Hull, annual in Jorion, semiannual in Tuckman) yet older practice problems tended to neglect explicit treatment of this assumption. We routinely sent that feedback to GARP, as perhaps others did too. GARP was responsive, and the modern practice problems are much clearer about this assumption.<sup>2</sup>

On behalf of my customers and myself, this memo has three goals. **One**, to warn you of reputational risks that are accruing to the FRM product, at least online. **Two**, to convince you there is an opportunity for improvement in the FRM product. **Three**, if those arguments are convincing, to learn that the Committee will take *some concrete steps* to make improvements. With that agenda, at the end of this memo, after asking important questions, I offer recommendations and suggestions.

---

<sup>1</sup> This memo defines “FRM syllabus” as the Study Guide (readings) plus Learning Objectives (LOs), and “FRM product” as the candidate’s FRM experience, including the syllabus, practice exams and the exam itself.

<sup>2</sup> Our contributions include, for example, relaying a need for consistent definitions (e.g., ratio-consistent information ratio), consistent calculations (e.g., surplus as risk), realistic quote conventions (e.g., forex quotes), many suggestions to convert “calculate” LO verbs to less demanding “describe” or “identify” action verbs, and many corrections to the source material, often submitted directly to the author/publisher; e.g., Malz, Meissner, Miller, De Laurentis.

I am in a good position to speak to the *online* reputation of the FRM product. Our training is provided entirely online. Some of our forum boards are freely available to public registrants (aka, members), such that customers of competitor EPPs are welcome to get help, which some do. For example, as Kaplan Schweser doesn't yet host an *always-on* support forum, we routinely help some of their customers, which we are happy to do; we perceive Kaplan to be professional, and their FRM product has a good reputation. Our forum has over 50,000 members, more than 9,000 discussions, and I personally have posted more than 15,000 times (where my positive approval rating is 99.63%). Most of my public- and semi- public forum posts are *replies* to FRM-related or finance subject-matter questions, or technical help in working out practice questions (ours or somebody else's). I also give limited support on our YouTube channel that features financial topics; last *month* this channel experienced a watch time of 443,000 minutes (7,400 hours viewed).

We host the world's most active public/private online forum devoted to FRM preparation<sup>3</sup>. As a longtime EPP, the issues itemized below are not new to me. What *is new* is social media, its reach and its velocity. When I started teaching the FRM in 2008, there was neither much online conversation about the FRM syllabus, nor much available feedback after the exam. Today, online feedback is swift and plentiful, both before and after the exam.<sup>4</sup> While writing this memo, I can clarify (or confirm) some of the cited examples with my online "trust network," almost immediately. It's great when you need help! On the other hand, social media has created customers who expect timely responses (sometimes unreasonably: a few customers have asked for always-on live chat!). Especially in the two months prior to the exam, subject-matter questions outpace our ability to answer, and we rely on our best customers to help other customers. Such is the double-edged sword of today's all-too-convenient online landscape.

My criticisms below do not mean to subtract from the admirable job GARP performs in *marketing* the FRM product and growing its customer base. Further, I am grateful to Kristina Jaoude and Bill May for responding to my concerns over the years. This memo certainly does not mean to criticize any individual. I enjoy teaching the risk curriculum. My commitment is to my customers, their experience, our business (which correlates with the continuing growth of the product), and the long-term health of the FRM.

I wish to remind us why high standards are important: high standards are important because FRM candidates are customers. Many companies are in the convenience business: they create customers by saving their customers' time, lowering their cost, or otherwise by making their customers' lives a bit easier or better. Not this product! The FRM product is the ultimate *inconvenience*. It demands a sacrifice from its own customers. Last week one of our customers, who is obviously studying in the evenings after work and revising problems on the weekend, lamented his "lack of quality time spent with my kids while they're still young enough to think I'm cool." A few lucky candidates can glide through easily. But the distribution of effort is quite disperse. Every year we host a few candidates who are visibly and genuinely "all in." This year we have a Part 1 candidate who is revising every single end-of-chapter question among assigned Hull chapters, in sequence; she has easily spent *several hundred hours*. For the outlier customers, their effort is basically heroic.

---

<sup>3</sup> Our public/private forum is located at <https://www.bionicturtle.com/forum>. The upper boards are entirely free to site members who register; these boards include FRM Topic boards (e.g., P1.T2. Quantitative Methods). The lower boards require a paid membership and primarily contain our database of >4,000 original practice questions which are "socialized" via thread discussions.

<sup>4</sup> As Kristina and Bill know, during exam weekend we must nowadays deliberately manage a "blackout" of candidate feedback because last year, some candidates posted exam information before the exam was concluded in all locations.

Of its typical customer, then, what does the FRM product expect?

- \$750.00 to \$1,050.00 for the Part 1 exam, *plus* \$350.00 to \$650.00 for the Part 2 exam. These fees do not include the purchase of study materials. Total fees can range from \$1,100 (early registrations) to \$1,700 (late registration). After including the purchase of study materials, total cost can exceed \$2,000.00. (apparently up to 30% register but do not sit for the exam. Probably some percentage of these “non-sitters” become overwhelmed by the syllabus. GARP offers no refunds.)
- An average of 240 hours per level (according to the GARP website), but the actual time invested by customers exhibits a widely reported range. Many candidates report studying 300 to 400 hours per Part, and some candidates report spending more than 400 (or even 500) hours for each Part.

### **GARP's official Practice Exams exhibit a high error rate and consequently send a poor signal**

The questions used in GARP's current 2018 Practice Exams are largely recycled from the prior three years (2015 - 2017). Many errors in this set have already been identified and corrected. We at bionicturtle.com contributed, by my estimate, at least twenty material corrections. We take little credit for *noticing* mistakes in the Practice Exams: >90% of our corrections are originated by customers (or forum visitors) who first notice a problem and then ask for my opinion, or our community's input, to help resolve their doubts. The typical candidate is rarely confident, often surprisingly unconfident, in their own isolated view. As GARP's practice questions are (on average) *less difficult* than ours, in most cases I can quickly render a decisive opinion. In some cases, I await collaboration or explicitly enlist the opinion of one or more members who has an expert reputation. GARP does not issue a public changelog (although we did receive one errata for the 2017 Practice Exam), so we cannot speak accurately to our share of the contributed corrections. We do observe that our visitors do not always utilize or possess *exclusively* the most recent Practice Exam. In some cases, somebody has told them that GARP's recent practice exams are essentially similar, such that only one version is needed; in other cases, they collect several years of practice exams with the goal of collecting as many “official” practice questions as possible. With respect to GARP's practice exams, by far the most common question we get is "How far back should I go when revising GARP's Practice Exam?"<sup>5</sup>

To my knowledge, GARP does not publicly *date-stamp* or *version-stamp* their Practice Exams; e.g., ver\_1, ver\_2, ver\_05\_07\_18. Consequently, especially lately, some visitors do not realize they are viewing a different version of an ostensibly identical question. Our orientation with respect to the Practice Exams is reactive. Due to their persistently low quality, we do not promote their use, but we also do not discourage their use. I often remind my audience, by illustration, that “even a flawed question can be instructive.”

For working purposes, I want to define a metric: **Original Error Rate (OER)** is the rate of error among the set of *original (aka, initial) versions* of individual practice questions in the last five to seven years. My definition of OER *includes* material imprecisions but *excludes* immaterial imprecisions (Immaterial imprecisions occur frequently in all questions, even among the best authors and editors. This is when a question is strictly correct, but the stress-test of actual usage reveals an opportunity for noticeable improvement).

---

<sup>5</sup> Historically our standard reply to the question, “How far back should we go?”, was: about five years. Last year, I backed off this guideline because the errors were creating too many support problems and extra work for us. Now, I try to support only the 2018 Practice Exam; and we recently are trying to direct more questions back to GARP.

Based on our documented sample, I am confident the **Original Error Rate (OER) on GARP's Practice Exams exceeds 30.0%**. This is my conservative estimate; as new errors are discovered (most exam candidates defer utilization of the Practice Exams until the final month or so before the exam), I believe the OER probably exceeds 35.0%, and incredibly, may be nearer to 40.0%.

Most of our forum visitors do not fully appreciate this reality: the majority of FRM Candidates are uniquely preoccupied by the daunting challenge of learning new, difficult material (also, most of *these* posts are not displayed in the upper, free boards). But, among our best customers, this has been noticed. In the last two years, criticisms have increased in frequency and intensity. A high defect rate in the official Practice Exam(s) has two effects:

- It engenders confusion in new learners. This aspect may be hard to appreciate for those without daily (or weekly) customer contact. New learners tend to presuppose GARP's official materials are accurate. I have spent countless hours helping some customers realize that GARP has made a mistake in a calculation or concept. Invariably, the follow-up question is, "Okay, so how will they ask about this on the exam?"
- It erodes confidence in the quality and fairness of the actual exam. *Obviously*. It has eroded my own confidence in the exam.<sup>6</sup>

If such an error rate is hard to believe, I empathize. Yet, as a *predictor* of the quality of the actual exam, there may be another worry. On average, the questions in GARP's Practice Exams operate at the **moderate** (*not advanced*) difficulty level. Specifically,

- The full range of Learning Objectives (LOs) include some decidedly *advanced* concepts. For many of these advanced learning objectives, there exist no examples of practice questions written by GARP. Given the error rate at moderate difficulty levels, it is natural to wonder whether GARP employs in-house expertise capable of consistently writing robust questions for advanced learning objectives.
- One of the **most prevalent complaints** posted after the exam is the observation that GARP's Practice Exams do not accurately represent the difficulty of the exam.<sup>7</sup> Many candidates are visibly surprised by the apparent shift in difficulty. This logically leads me to wonder if some of the perceived-to-be-difficult questions were simply bad questions. We know from hundreds of examples in our forum that, from the perspective of new learners, *a bad question often presents as a difficult question*. Many, many times in our forum this has happened: a visitor posts a question (from whatever source) because it is "too difficult", but analysis proves the question to be flawed. Some of our most time-consuming support threads relate to flawed questions that have led the learner down a blind alley.

---

<sup>6</sup> For flavor, here is a typical comment from a candidate who had helpfully identified and posted *several* errors over *several* days and who was understandably exasperated: "It is very annoying to see [GARP] constantly making mistakes like these...and [it] also costs me tons of time to check, reference back to books ... I am wondering if they will make mistakes on the real exams, that will be terrible." (<https://trtl.bz/2HMCH0c>). **To maintain our forum's integrity**, including our own transparency and accountability, we do not generally moderate (edit) member/subscriber posts, except for spam. In this way, customers also very *frequently* identify mistakes in our materials (we also have forum threads dedicated to our own errors).

<sup>7</sup> The follow example comments are *typical* of this *extremely* common feedback point: "The GARP practice exams are definitely not an indication of the level or type of questions given"; "The level of difficulty was not even close to GARP 2016 practice exam."; "The exam was a lot tougher than the GARP practice exam."; "The GARP practice exam was not representative of the actual exam. Actually, I think it was misleading in the sense that it made one underestimate the topics being tested, the level of difficulty and complexity"; "I felt all my practice exam scores were thrown out the window. Actual exam had lot more complicated questions with very long calculations."

To illustrate for the reader, below I have summarized brief feedback on the *first ten (10) questions* from the 2017 Part 2 Practice Exam. This is a very *partial* sample of our database. Please note these questions had *already been edited* (recycled) from the prior year such that, given an entire year, it was reasonable to expect a fully corrected paper. Among this small sample, the error rate is 50.0%, which is above the general average, but their nature is representative of mistakes made throughout the history of all practice exams.

1. Correct EVT question (difficult calculation applies Dowd's EVT POT formula)
2. Correct normal versus lognormal VaR question (tedious calculations)
3. **[Glaring Errors]** The original asks for default probability using the KMV Model but (1) utilizes a default threshold formula unavailable to current candidates; i.e., from discontinued De Servigny Chapter 3, and (2) incorrectly applies  $PD = N(-DD)$  by assuming a Merton model. This was a baffling confusion, noticed by candidates, between the KMV and Merton distinction, which has been *explicitly taught for a decade* in de Servigny; e.g., 2009 LO = "Compare the Moody's-KMV's equity model with the Merton model."
4. I defended this question as basically **correct** (it repeats in 2018 without further edit) because the correct answer is *clearly* the best answer. Some customers argue that parametric distributions can include non-normal skew; we tagged this question as "can be improved."
5. Not an error, but most candidates have trouble applying the referenced Malz, which requires skill to follow<sup>8</sup>. For these concepts, we prefer to refer candidates to the more accessible Hull or even Saunders, who are no longer assigned (but upon which historical examples are anyhow based).
6. **[Confusing Error]** The original contained a typical example of GARP's tendency to confuse various *default probability* definitions over the decade. The question supplies "marginal" default probabilities, but the reference (Malz Chapter 7) defines *marginal* PD as the first derivative of the continuous cumulative PD,  $\lambda * \exp(-\lambda t)$ , a metric unfamiliar to most candidates<sup>9</sup>. GARP meant to supply "conditional" default probabilities. GARP's published confusion on the flavors of default probability engenders at least three queries per month on our forum. Other instances of this theme include: Question 29 incorrectly confused *unconditional* default probability with *conditional* default probability; Question 10 (Practice Exam 2015), Questions 33 and 76 (2016 Practice Exam) incorrectly distinguished between *joint* versus *conditional* versus *default* probabilities. Question 78 (2017 Practice Exam) contained a flaw, noticed by no less than four candidates, and used a formula not supplied in the solution's referenced Gregory. On a positive note, the correction in 2018 produces a **very good** question, in my opinion!
7. Correct, classic, easy question about credit VaR. An example of a good question.
8. **[Error]** Incorrect pension surplus solution discovered by candidates (reported to GARP, who fixed and re-issued)
9. **[Error]** Incorrect VaR solution reported by candidates (reported to GARP, who fixed and re-issued)
10. **[Multiple Errors]** RAROC calculation contained *at least* two errors, all noticed by candidates, including inconsistent numerator/denominator, which is the first rule we teach (reported to GARP, who fixed and re-issued). Question 41 also contained a mistaken RAROC calculation<sup>10</sup>.

Our forum experienced visible, sometimes vocal, frustration with the quality of GARP's 2017 Practice Exam. I voiced these concerns to Kristina and Bill.

---

<sup>8</sup> To appreciate the challenge of applying Malz, we recommend reading Malz for these learning objectives: Calculate risk-neutral default rates from spreads. Explain how a CDS spread can be used to derive a hazard rate curve.

<sup>9</sup> As one candidate posted, "it seems we never ever come across Malz' definition of marginal default probability."

<sup>10</sup> A fuller explanation of this RAROC is here <https://trtl.bz/2HHXMnv>

I have not proactively analyzed the full 2018 Practice Exams (we have our own database of 4,000+ original practice questions to support). However, I have responded to some customer questions, as part of my daily customer support ritual. My customers have observed *three* additional errors in the last three weeks:

- **Part 1 Question 30** contains a fundamental misapplication of central limit theory (CLT). This error has been reported several times, apparently going back to 2014. This sends a *terrible* signal because neither the author(s) nor editor(s) seem to understand basic statistics.
- **Part 1 Question 13** on VaR mapping. This is a subtle but revealing mistake. The VaR mapping syllabus assignment (Jorion Chapter 11) has been unchanged for at least a decade. The authors (and editor) of Question 13 have written a credible question for somebody who has read the chapter superficially because it accurately queries the features of Jorion’s *literal* scenario. Ironically, less knowledgeable readers are likely to choose the correct answer, but more informed readers might pause. Two of our customers, who are engaged with the topics at meaningful level, noticed it contains a flaw<sup>11</sup> I did not notice, I cannot take credit! My own questions tend not to commit this sort of flaw because I *construct* the assumption scenario to ensure coherence. I wish this were an isolated example, but it fits a pattern. For many of the advanced “knowledge points” assigned to customers, where we all expect GARP’s authors to exhibit expertise (certainly on a reading with over a decade of seasoning!), the practice questions belie this expectation. In this situation, the subsequent *follow-up question is almost always* some variation on the question: “Okay I get how they did it here, and now I get how it should be done, but what should I expect them to put on the exam?”
- The third problem (so far) revealed in the 2018 Part 1 Practice Exam was discovered on 5/14/18 by a candidate; like other errors, I remind the reader the 2018 version was already the beneficiary of several edits over multiple years. **Part 1 Question 33** is a disturbing artifact because the authors (and reviewers) of this question ***do not understand the cost of carry model*** that has been taught to candidates—delegated to the same John Hull material and essentially unchanged—for about a decade. GARP’s authors seem to believe that the cost of carry model generates an expected future spot price per *three* (3) references to the “future value of the spot price.” The explanation includes:

“Future value of the spot price =  $S_0 * \text{EXP}[(\text{risk free rate} - \text{dividend yield}) * 1.25] = \text{USD } 3,763.52$ ”

But this is wrong. It’s important to understand why GARP’s shallow confusion matters. Hull’s great textbook often isn’t directly about risk, it’s firstly about valuation (aka, pricing); e.g., the swaps chapter concludes with swap *valuation*; while the *risk* of swaps is elsewhere and builds on the pricing foundation. Hull’s *cost of carry* model (also covered in McDonald) concerns the theoretical price of the futures contract,  $F(0)$ . Risk enters the model, in one way, by virtue of the *theory of normal backwardation*: the hedger compensates the speculator, who is long the futures contract such that  $F(0) < E[S(t)]$ . This is the expected relationship when the commodity is an equity index and, therefore, has positive beta. In Question 33, the expected “future value of the spot price” is *greater than* 3,763.52 by a *multiple* of  $\text{exp}[(k-0.050)*15/12]$ ; i.e., the short hedger has an expected loss on the naked derivatives trade. From the perspective of *risk*, this is a nuance of the model that matters to risk learners. This is a concept some of our Part 1 customers already know and has been stable in the syllabus for over a decade.

---

<sup>11</sup> If you can’t figure out the flaw, here is a hint: what is the defining feature of the Vasicek Model in Tuckman’s Chapter 9 (The Art of Term Structure Models: Drift)?

## Evidence suggests actual exams have not been sufficiently robust and fair

I am unfortunately confident that an audit of past actual exams would reveal a significant number of material errors. I believe a subset of these mistakes might have the perverse effect of biasing against candidates who are better prepared. Further, my *mere hypothesis* is that some questions may have been flawed yet were simply *perceived* as difficult. We have directly observed how some errors tend to *trigger a pause in our more careful and engaged students* (aka, customers). My evidence includes:

1. In selling the Practice Exams, GARP claims they are "based on a sample of questions from previous FRM [Part I or Part II] Exams ... " If this advertised statement is indeed true, then I can claim with *high confidence* that at least some **actual FRM exams** contain multiple material errors. Although the current 2018 Practice Exams enjoy the benefit of years of submitted corrections and third-party assistance, to the extent these questions were *sampled from actual exams*, obviously they appeared in their *original versions*, not today's edited versions. Statistically, this is a large sample ( $n > 30$ ), so conditional on a random sample, a confidence interval for the true error rate can be estimated.
2. GARP was forced to admit a mistake on the actual November Part 1 exam<sup>12</sup>. Although we observe reports of suspicious questions after every exam, the lack of transparency (and restrictions on candidates) has precluded any reliable verification. However, this question was *uniquely brief and simple* such that—unlike with elaborate or difficult questions where ex post recollections vary—the recollection was unanimous. This mistake was surprising because the solution misapplied a basic unexpected loss (UL) formula that has been **stable in the syllabus for over a decade**. I do not know who and how many experts reviewed this question, but every single person who reviewed this question *failed* a basic test.
3. There exist conceptual subtopics where GARP's expert authors (at least to the extent they write, edit, and/or review practice exam questions) reveal a lack of depth. A good example is *Portfolio Value at Risk (VaR)*. This topic should be deeply familiar to everybody engaged: the exact same reading (Jorion Chapter 7) has been assigned for over a decade. It has never been updated. But for most of the last decade, the associated practice questions supplied a *mathematically impossible* assumption set.<sup>13</sup> As recently as 2017, Question 60 repeats the pattern. As far back as 2009, candidates have been expected to "Define, compute, and explain the uses of marginal VaR, incremental VaR, and component VaR." But it is not obvious that GARP has mastery over this learning objective, much less the entire chapter. In case the reader thinks our standard is too high, consider the reason this first came to my attention. Each of the above metrics can be approached in (at least) two ways. A robust assumption set will return the same solution regardless of the approach. This first came to my attention because a knowledgeable candidate solved for marginal/component VaR using a valid approach, but his solution was not a given choice. I think it's a reasonable hypothesis that, on the actual exam, most of these portfolio VaR questions over the last decade suffered this problem.

---

<sup>12</sup> Here is my comment about the error <https://trtl.bz/2IqrTj7>

<sup>13</sup> For over eight years, GARP's expert authors and reviewers apparently failed to realize the position betas are "with respect to the portfolio (which contains the position)" with a necessary relationship to the other metrics.

## The syllabus may not be receiving enough attention

Hoping that I have earned the reader's attention, I shift to more constructive feedback; i.e., because the following issues can be fixed with *better process design*. Our forum has accumulated thousands of posts that collectively constitute *unstructured* feedback on the **details** of the FRM syllabus. Some feedback is more subjective than objective.<sup>14</sup> Further, my points below embed *my own interpretations and opinions*, albeit informed by customer engagements. While organizing this feedback, I decided somewhat arbitrarily to arrange the feedback into three categories: Recency, Coherence, and Relevance.

### I. Recency

Recency refers to readings that seem to remain in the syllabus due to inertia. We do not understand why some of them haven't been updated, much less refreshed, in over a decade. For example:

- **Linda Allen's Chapters 1 and 2** (the *first* two readings in Topic 4) were published in 2004 (**14 years old**). Every year, Part 1 candidates spend hours grappling with Allen's technical hybrid VaR interpolation. It's subtle enough that that we've helped at least one competitor EPP fix their misinterpretation of it (it is often misunderstood as an *interpolation of* interpolated loss quantiles). I don't expect it will ever be tested on the exam. Not unrelated, if a candidate wants to "Explain structured Monte Carlo ..." (per the LO), this reading is superficial and outdated to the task.
- **Kevin Dowd's Chapters 3 and 4** (the *first* two readings in Topic 5) were published in 2005 (**13 years old**). The more advanced sub-concepts seem only remotely testable. There is an LO that expects candidates to know how to "estimate risk measures by estimating quantiles" but GARP's initial attempt contained a methodological mistake noticed by one our best customers, which he posted in a batch of *seven errors that he captured in his first attempt* at GARP's 2016 Practice Paper. And this is a *relatively* easy section in the chapter. Every year we help some ambitious candidates understand non-parametric density estimation (NDE) per the LO "Describe historical simulation using non-parametric density estimation," or for that matter the mathematical depths of general risk measure, coherence and expected shortfall (ES). But after 10+ years of this reading, I cannot speak to GARP's in-house mastery over the tougher Dowd sections, much less whether they might be tested.
- **De Servigny Chapter 2** (External and Internal Ratings) was published in 2004 (**14 years old**). This chapter predates the crisis. It sincerely makes me a little sad to embed this utterly outdated material (e.g., empirical assertions based on data from the 1990s) into our own product; I feel like it reflects a willingness to waste their time.
- **Jorion's Value at Risk (3rd edition)** was published in 2007 (**11 years old**). Portfolio Risk (Chapter 7) features the previously discussed Portfolio VaR, which clearly GARP has struggled to master. Every year candidates study the *latter section* (i.e., risk-minimizing position, portfolio management) but I cannot speak to GARP's mastery of them, or whether they will be tested.
- Dowd's **Extreme value theory (EVT)** chapter was published in 2005 (**13 years old**). The first question in the current Part 2 Practice Exam tests this reading, I encourage the reader to see if they can solve it. I'm interested in this topic; perhaps something new has developed in 14 years.

---

<sup>14</sup> And some feedback should be ignored. Among the best pieces of advice I got when launching an online community was from Brian Clark who told me, "Ignore individual feedback." He meant that it is unwise to take seriously every single comment; rather, you should consider reliable only feedback that accumulates toward a consensus and *emerges as a pattern*. We don't really count something as feedback until at least three visitors post.



- Litterman's Chapter 17 (Risk Monitoring and Performance Measurement) **was published in 2003 (15 years old)**. I am not aware of any practice questions that refer to this reading. It only appears on our forum when candidates are trying to decipher/reconcile its terms (e.g., risk management unit, alpha, tracking error) with other readings. One of its ten LOs is “Identify sources of risk consciousness within an organization.” (The answer is: banks that lend, boards, and investors). The dashboard examples are possibly from another era.
- Stulz Chapter 18 (**Credit risk**) **was published in 2002 (16 years old)**. Every year some candidates struggle mightily with section 18.1 especially. In recent years, some of us have dismissed some of its findings in favor of updated material<sup>15</sup>. When expected to “Compare and contrast different approaches to credit risk modeling, such as those related to the Merton model, CreditRisk+, CreditMetrics, and the KMV model.,” candidates are assigned this 16-year old reading.

## II. Coherence (consistency)

Coherence refers to the loose coordination among many syllabus elements, either: conceptual, semantic, definitional or symbolic/notational. Much of this is an *inevitable* consequence of an **anthology-based** approach coupled with little or no editorial “glue” to connect disparate authors. You’d be surprised at how many discrepancies—ranging from tedious notational differences to occasionally thought-provoking enigmas—are identified by candidates.

Below are some examples culled from our forum discussions. As inconsistencies are far too numerous to list in their entirety, below I culled a few examples raised by candidates in just the last two weeks.

- **Absolute versus relative VaR**. This is a perennial confusion for candidates. VaR is just a quantile that can be applied in each of the major risk categories (aka, MVaR, CVaR, OpVaR). For example, CVaR can be relative (UL) or absolute (EL + UL) and the default can vary by author; e.g., Malz versus Basel. We teach the absolute/relative distinction and insist it should be explicitly defined, but to my knowledge, GARP doesn't employ this distinction. More than once, candidates ask whether it is true that  $CVaR = UL(\alpha) - EL$ . Probably they ask because the 2015 Practice Paper explains that “The 95% credit VaR corresponds to the unexpected loss at the 95th percentile minus the expected loss.” I do not think  $UL - EL$  is meaningful in any context (this mistake was reported/corrected).<sup>16</sup>
- **Surplus at risk (SaR)** has only ever been assigned under Jorion’s Chapter 17 (VaR and Risk Budgeting in Investment Management, 2007). The historical approach is uneven. We teach each of three (or four) valid approaches, but I’m sympathetic to an EPP competitor’s claim that it “won’t be tested.”<sup>17</sup> because GARP has given no reliable guidance here.

---

<sup>15</sup> This recent post illustrates a typical stumbling block in Stulz 18.1 <https://trtl.bz/2vWzDEB>

<sup>16</sup> For example, one customer posted: “GARP’s definition: The 95% credit VaR corresponds to the unexpected loss at the 95th percentile minus the expected loss...” When I **explained this is incorrect**, he replied “Thank you so much. Sorry I couldn't read your response earlier, but I've been getting my \_\_\_\_ kicked (I'll let you fill in the blank) in the practice exams so I've been spending some time on that. I do have another question that's somewhat related. When GARP gets an error pointed out to them, do you know if they have a process for going back and reconciling with old exams to see if similar question was present there? I know it may sound like a dumb question but given how many errors I (and others) have seen, I wouldn't be surprised if similar issues surface in the actual exam. On another note, wouldn't it be nice if someone gets an email saying they actually passed because GARP corrected a (or few) problematic questions.” <https://trtl.bz/2IsbGtx>

<sup>17</sup> See <https://www.analystforum.com/forums/frm-forum/91361965>

- De Laurentis' **Key Concepts of Credit Risk** is my least favorite new reading in the syllabus, partly because I care about the underlying concepts (e.g., logit models) that are so compelling. But not in this reading. Apparently, the source still contains multiple errors; incorrect formulas enabled by the lack of any illustrative examples. The Harvey balls do not edify. The Chapter 3 syllabus assigns *fully 14 learning objectives* in high-density fashion (e.g., heuristic and numerical approaches), yet there are no practice exam questions and no end-of-chapter questions. We have no clue how they might be queried, or if they will be queried. For members who want to learn something job-useful about credit risk, we are sending them to some great courses at [datacamp.com](https://www.datacamp.com).<sup>18</sup>
- **Liquidity cost** and liquidity-adjusted VaR (aka, LVaR). Inconsistencies and/or mistakes in the practice exams engender confusion. For example, Question 2017.P2.35 contains at least two mistakes<sup>19</sup>
- **Default probability** definitions and various approaches, including variations on the legacy no-arbitrage approach<sup>20</sup>
- **Credit value adjustment (CVA) and xVa**: With respect to CVA, GARP's rotation from Canabarro to Gregory's three editions has plagued us with support work; e.g., Gregory's 3<sup>rd</sup> edition (xVa) changes key formulas, but it seems unlikely GARP is aware of these details
- **Other examples** where there is at least some degree of lack of coherence in the syllabus: VaR mapping, alpha, tracking error, and RAROC.<sup>21</sup>

---

<sup>18</sup> "Credit Risk Modelling in R" is a [datacamp.com](https://www.datacamp.com/courses/introduction-to-credit-risk-modeling-in-r) course that teaches logistic regression in a hands-on format <https://www.datacamp.com/courses/introduction-to-credit-risk-modeling-in-r>

<sup>19</sup> About this question, a customer wrote: "a huge waste of time to try figuring out what they [GARP] meant." <https://trtl.bz/2It9BgU>

<sup>20</sup> Here is a typical example of a candidate's confusion in regard to the no-arbitrage approach: <https://trtl.bz/2IsPfEE>

<sup>21</sup> Tracking error is queried dozens of times each year due to inconsistent definitions; e.g., in assigned Chapter 7, Jorion apparently defines active return as "tracking error" and tracking error as "tracking error volatility." GARP's inconsistent approaches to RAROC engender perennial confusion; e.g., while editing this memo today (May 10<sup>th</sup>) a candidate *this morning* posted "I just wanted to clarify the RAROC formula. I looked at both past [GARP] papers from 2016 and 2018. according to 2016 Q10, RAROC formula excludes return on economic capital from tax calculations, whereas Q41 from 2018 includes return on EC in the tax calculations. Which is the correct approach? I believe we should tax all revenues/income/costs/fees etc." <https://trtl.bz/2KbpxBW>

### III. Relevance

Relevance is an umbrella that seeks to cure three defects: *Discontinuous Assignments* are widowed/orphaned readings that do not stand alone very well. Diligent pursuit of the LOs tends to require supplemental study. *Potential Over-assignment* refers to patches of the syllabus where strict adherence to the learning objectives can be *extremely time-consuming*, but the payoff is unclear. *Frenetic Current Issues* refers to implications of swapping the *entire set* of Topic 9 readings *every year*.

***Discontinuous Assignments*** require supplemental effort. For example,

- **Grinold’s Chapter 14** (Portfolio Construction, published in 2000) is almost comically absent something educators call *scaffolding*. If true mastery is really wanted, candidates need to backfill their knowledge with preceding chapters. One customer wrote, “Sorry for the REALLY long question. The choice of GARP to include this chapter is really bad. You need a LOT of background to even understand what is going on and to understand what the variables represent.”<sup>22</sup> Yet we’ve observed that GARP only asks superficial questions about this material; e.g., Part 2 Question 63 is almost famous for puzzling our visitors.<sup>23</sup>
- **Diebold Chapter 5** (published in 2006) begins a four-chapter sequence in time series and forecasting. It’s hard to calibrate how seriously to take these learning objectives. On the one hand, to be able to truly respond to every learning objective requires diligent, time-consuming study. Diebold’s own end-of-chapter (EOC) questions are very demanding. On the other hand, the single 2018 Practice Question that refers to Diebold is a dead simple *conceptual* (not even quantitative) question about covariance stationary. Serious study requires EViews, the software used in the book.

***Potential Over-assignments*** refer to subtopics (or patches of learning outcomes) that seem unlikely to appear on the exam yet can be *heavy time-traps* for especially diligent candidates. This is a minor pattern among the post-exam feedback items: highly prepared candidates are prone to voicing frustration that they spent many hours studying for a section that was completely untested. Possible examples may include (this is a very partial list from my perspective):

- Calculus in Miller
- John Hull’s Interest Rate Futures chapter; e.g., many candidates spend many hours understanding the details of this chapter, including “Explain and **calculate** a US Treasury bond futures contract conversion factor.” (I would be surprised if any reader of this memo recalls how to *calculate* a T-bond conversion factor), and “Calculate the theoretical futures price for a Treasury bond futures contract”
- Tuckman
  - Risk-neutral derivation
  - Certain details of term structure models

---

<sup>22</sup> See <https://trtl.bz/2Irg8cf>

<sup>23</sup> The question applies Grinold’s  $\sigma(\alpha) = \text{volatility} \times \text{IC}$ . We infer a superficial application because the original question contained an error, as several customers noted. The 2018 fix has prompted a new question from Karim (one of our best customers): “why are there [now] fewer alphas than stocks?” The back-and-forth on this question, in service of correcting GARP’s error and discerning the question’s intention, in my opinion, is not very productive: there are much better questions that could be written to *really* test a meaningful understanding of Grinold.

- Jon Gregory: GARP has densely assigned learning objectives (LOs) to several chapters. But the practice exam questions are limited to a few, easier learning objectives (LOs). We have no evidence to suggest the harder LOs will be tested. Related, the *current* Practice Exam refers throughout to the text's prior edition (Counterparty Credit Risk and CVA, 2<sup>nd</sup> Edition). But the syllabus assigns the recent edition (The xVA Challenge, 3rd Edition). I have no idea if GARP is aware of Gregory's nuanced changes between editions; e.g., we are forced to support the related confusions.
- Basel: I could write an entire memo on Basel alone. In fairness, the Basel regulations are inherently detailed and dynamic. Still, if GARP's motivations do include "we care about candidates' time," there is room to provide better prospective clarity to candidates. Just one small example: this week a candidate posted "Looking at question 10 for the 2018 sample paper it looks like we are expected to memorize the BI buckets!!! Unbelievable!"<sup>24</sup> (Legitimate inference. The associated LO is: Calculate the operational risk capital requirement for a bank using the SMA)

***Frenetic Current Issues (Topic 9):*** Every year GARP **replaces the entire set** of Current Issues readings. Yet I think the Practice Exams do not contain even a single question that references *Current Issues* in at least four years. Why should we be confident the actual exam will contain approximately eight high-quality questions (i.e., 10% weight \* 80)? Without examples, EPPs cannot shape an expectation about the sort of questions that might be asked.

I wonder if the inability to write Current Issues practice questions is symptomatic of too much turnover. Maybe GARP shouldn't rotate readings faster than it can write (at least one or two or three) practice questions. If GARP holds itself to the same standard as its customers, maybe readings should not change faster than new questions—at least some new questions—can be written.

More importantly, does it help candidates to change these readings (and therefore the entire set of learning objectives) every year *in their entirety*? I personally thought 2017's virtual currency assignment was a terrific choice, and one of the best summaries of Bitcoin that I've read. It's gone. New for 2018 includes a strong reading on machine learning. My prediction is that machine learning will still be relevant in 2019, maybe even in 2020.

---

<sup>24</sup> See <https://trtl.bz/2IrEoLn>

## Questions for the Board and the FRM Committee (optional)

Below are my questions for the individuals with some responsibility or participation in the governance of the FRM product.

- *Excluding* purchased study materials (i.e., not including the books that cost \$300.00 per part), I estimate the FRM product generates revenue (enrollment plus exam fees) of *at least* \$30.0 million per year. Fees collected for registrants who do not sit (as no refunds are given) is probably, at least, a \$5.0 million annual business by itself. Question: Compared to profits, *is the organization spending (investing) sufficiently in order to meet its advertised promises--implicit and explicit—to customers?*
- GARP’s website reads, “What does the FRM Committee do? The FRM Committee plays a vital role in preparing the FRM Exam. The Exam topics and their respective weightings are reviewed yearly to ensure the FRM Exam is kept timely and relevant.”
  - But the *high-level* topics (e.g., Topic 2 Quantitative Methods) and weighting have not changed in several years. Does the FRM Committee have any role and/responsibility in the lower-level syllabus; i.e., assigned readings and their associated learning outcomes?
  - If not, what are the Committee’s *hands-on* responsibilities, if any?
- **Who at GARP is a Certified FRM?** (We can’t identify any Certified FRMs among GARP’s team)
- What is the philosophy, if any, with respect to decisions to *change or update* syllabus components?
  - Why is EPP input not considered? E.g., effort required to create new material for new readings.
  - Why did the syllabus rotate *furiously* for most of the past decade, churning 20% or 30% or more (e.g., econometrics rotated rapidly *without* key conceptual shifts but primarily workload implications for EPPs and candidates) then suddenly stabilize in 2018 (except for the usual rotation of all Current Issues)?
  - Is customer input considered at all; e.g., input from on-the-job FRMs?
- Why hasn’t GARP created **any original content** to create editorial “glue” (consistency) that could support the inherently discontinuous anthology-style syllabus? <sup>25</sup>
- Regarding the Practice Exams, why is GARP’s capability seemingly limited to the generation of a single set of Practice Exams, *which itself required three or four years of iteration* in order to effectively proof (and apparently, with much external assistance)?
  - Why hasn’t GARP generated practice questions for many advanced sections of the syllabus, content which is explicitly assigned to customers via the learning objectives?
- Is it reasonable to expect GARP to employ in-house (or hands-on) expertise on all syllabus topics?
  - If so, who are the resident expert(s)?
- Is it reasonable to expect that GARP should be able to provide candidates and EPPs with better *prospective clarity* in terms of exam expectations; i.e., to better coordinate the relationship between learning objectives (which, after all, constitute the realistic guidance), the implications on customers’ precious time (which varies by assignments), and the actual exam expectations?
- With respect to the actual exam questions:
  - Who exactly writes them?
  - Who exactly reviews them?

---

<sup>25</sup> This lack of “glue” effectively *transfers workload* to EPPs like ourselves; further, it’s conceivable this vacuum “dulls the senses” and encourages the sort of detachment that allows for the writing of practice questions that effectively render candidates and EPPs into beta testers. Given the lack of leadership concerning syllabus inconsistencies, I have started an open-source **glossary** which will be shared with other EPPs. I am doing this because there are perennially so many questions about perceived (or real) terminology differences among syllabus authors.

## Some Recommendations and Suggestions

- Benchmark the FRM product's cost of goods sold (COGS). Increase product spend accordingly to achieve threshold product quality according to commercial, professional, and ethical obligations.
- Clarify the *governance responsibilities* with respect to syllabus quality, practice exam quality, and the linkage between the syllabus, the practice exams and the actual exam.
- Clarify the linkage between business model, customer mission, role of EPPs and syllabus. Including:
  - Add an explicit reverse-engineering guideline to the philosophy, such as: *Let us deeply respect the time invested by our customers (and our EPPs).*
  - Have deliberate intention (and consideration for stakeholders) in the *pace of change in the syllabus*. For example, perhaps the ideal syllabus strikes a balance between continuity and innovation. Avoid changing readings for the sake of change.
  - Formalize feedback loop(s) to promote continuous improvement; e.g., the promotion of a consistent foundation of durable *career-enhancing* skills.
- Introduce accountability--and some degree of transparency--into the currently opaque exam process.
  - **Submit the May 2018 exam to an expert, third-party audit** <sup>26</sup>
- Require key FRM product team members to either **(i) hold the FRM designation or (ii) to immediately enroll in the FRM**. This will begin to cure the apparent detachment from the syllabus and encourage intimacy with the product sold to customers.
- Borrow selected *best practices* from the CFA Institute
  - To me, this seems like “low hanging fruit” and the easiest way to make improvements. The CFA Institute is already a role model with respect to methodology, quality, communication, and accountability. If GARP borrowed even a few practices, improvements might be noticeable.
- Introduce a formal liaison mechanism with exam prep providers (EPPs).
  - Use a project manager (PM) platform to track issues
  - Introduce and make available the in-house expert(s)
  - Make public (or semi-public) a *changelog* for the syllabus and practice questions
  - Consider hosting an online private forum for EPPs to disseminate key updates
- Conduct a hands-on, expert audit of the learning objectives. Specifically, prune and streamline objectives, especially where they are over-assigned. Check carefully against the assigned reading. <sup>27</sup>
- Identify selected domain areas where GARP can **demonstrate leadership by creating its own content**. (Not blanket coverage, sharpshooter-style)
- Plan for future *innovations*. Ideas (for example):
  - Survey on-the-job Certified FRMs to identify *forward-looking* career skills that will be essential
  - Prepare for the possibility of electronic exams; e.g., problems and case assisted by Excel (rather than calculators) which helps candidates gain a skill that is actually useful in the workplace
  - In Topic 8 (Investments), shift away from tired CAPM reread and inchoate Jorion sections in favor of modern, model-based applications with actionable utility to portfolio management
  - Give more weight to career-enabling technological *applications* that intersect with the modern risk manager's role; e.g., data science, fintech developments, dataviz (risk communication)

---

<sup>26</sup> We are not yet requesting a historical audit of actual exams (our data suggests the results might be catastrophic)

<sup>27</sup> Space precludes a full listing of the *numerous* opportunities to streamline (or render more precise) many learning objectives (LO). For example, John Hull eliminated the swaps valuation approach “based on two simultaneous bond positions,” in his 10<sup>th</sup> edition update, so that LOs can match by deleting “Calculate the value of a plain vanilla interest rate swap based on two simultaneous bond positions.” And tiny example of better precision: a Miller LO reads “... calculate joint probabilities using probability matrices,” but joint probabilities are *observed* as elements of a probability matrix; you might *calculate* a conditional probability given a joint probability matrix.

## Conclusion

Thank you for your time and attention. I hope you agree that my criticisms are justified and constructive. I've taken several hours—during the busy customer support season immediately before the exam—to collect and organize this feedback *partly* because I happen to believe these illustrated, low standards cannot, in the long run, survive the gradually rising tide of technology-enabled (e.g., social media) customer expectations. And it's not merely that social media is already *forcing* transparency. EduTech disruptions are coming fast. MOOCs like Coursera and Udacity already offer (micro-) credentials, at a much lower cost, with proven ability to bypass slower institutions and enable graduates into technology jobs; they are now visibly encroaching into other verticals. While we administer paper exams with hand calculators (and weeks- or months-long waits for experience verification), they enable superior peer-review at scale and analyze clickstream data with pattern recognition to identify and correct quiz problems *in real-time*. Our Current Issues has a reading on machine learning; they are using machine learning to *optimize* customer experience. GARP (and the FRM Committee?) added/deleted only two Credit Risk readings in the 2016-17 update; one was De Laurentis, which *more than a year later* still has fundamental errors and conflicts that are confusing candidates.

But my *main* motive is simply that I want the best possible product for my customers. Or, at least a better product. In my imagination, the FRM has the possibility of realizing its *potential* as the “gold standard,” aka, the highest standard. But before it can be the gold standard, it needs to meet the standard that it expects of its own customers. It's just not enough to curate readings and copy headers into learning objectives, then come back months later and write shallow and/or flawed questions as part of an opaque exam process. If GARP expects its paying customers to understand the marginal VaR calculation or VaR mapping or scaled alphas or five different interest rate term structure models or more than fifty (50+) learning objectives on counterparty credit risk, it should understand them too. At least well enough to (i) offer candidates and EPPs prospective clarity and (ii) administer a robust and fair exam. These are the *reasonable requirements* given the costs involved, in time and money.

I hope I've given you enough feedback to do something. Nicole and I are deeply invested in the success and continuing growth of the FRM product. We interact daily with FRM candidates, and we hope that our efforts push the brand forward.

As I've mentioned a few times, my examples and quotations are (i) a very *partial* sample of our total database and (ii) are all documented. If you would like to probe further, please feel free to contact my colleague Nicole ([nicole@bionicturtle.com](mailto:nicole@bionicturtle.com)) or myself ([davidh@bionicturtle.com](mailto:davidh@bionicturtle.com)) and we will give you free, unrestricted access to the paid, member-only boards on our forum. Please also feel free to follow-up with me directly by email or my office phone (424-206-9425). Thank you!

## Appendix A: Specific Questions (“Short List”) for Trustees and Committee

1. GARP advertises the for-sale practice exams are "based on a sample of questions from previous FRM [Part I or Part II] Exams ..." and we have documented an (Original) Error Rate (among a large sample) in excess of 30.0% among these exact questions.
  - a. Are you **willing to publish, or make available, the past actual exams** (so that we can ascertain the exact, actual error rate)?
2. With respect to actual exams:
  - a. Who specifically writes the questions?
  - b. Who specifically reviews the questions?
  - c. Why is there no transparency with respect to quality control? e.g., no public feedback concerning identified errors, such as changelog or error log
  - d. Who is the “**highly respected independent third party**” and what exactly is its role?
3. Why is the **lack (or paucity) of Certified FRMs** employed by GARP tolerated? This unwillingness to “eat their own dog food” might partially explain (i) why the sample of published questions exhibits such a *low standard*, (ii) why candidates and EPPs experience suboptimal *prospective clarity* with respect to the exam and (iii) why the syllabus doesn't have a *tighter link* to the exam.
4. With respect to the Study Guide (readings) and Learning Objectives?
  - a. Who has a hands-on role in their development; i.e., review and update?
  - b. Who specifically at GARP has deep-level expertise and/or mastery in the syllabus (readings and learning objectives)?
  - c. Who specifically has the expertise to field technical questions about the syllabus?
  - d. Who specifically has the capability to write exam questions for the Current Issues (which change every year)?
5. With respect to the budget?
  - a. Is a sufficient percentage of revenue (or profit) spent on Cost of Goods Sold (COGS) and/or Research and Development (R&D)?
  - b. Why hasn't GARP developed any of its own content to plug holes in the syllabus; e.g., extremely old assignments?



## Appendix B: My “Response to Bill May’s Response”

Bill,

Thank you for your response. I appreciate the time you’ve taken to outline certain processes and efforts underway. Many of your efforts are *clearly* in the right direction and deserve everybody’s full support. Here are my specific reactions:

- **Practice Exams (PEs):** Your memo tries to attribute reputational problems with GARP’s Practice Exams (PEs) to the 2017 Part II exam. This understates their problems. *Every single PE published by GARP has been plagued* by multiple, material errors. The 2015 PE was worse than the 2016 PE was worse than the 2017 PE; the 2017 PE was the beneficiary of cumulative corrections (many submitted by us). Candidates on our forum ***continue to identify conceptual mistakes in both Part 1 and Part 2 of the 2018 PE***. This cannot be chalked up to a snafu or version madness with respect to a single PE. As my memo shows, the serious problem is how they have consistently (i.e., every year) exhibited material or conceptual flaws that *signal a lack of mastery* of the material assigned to your customers.
- **Actual Exam Questions:** The PE exams exhibit a documented original error rate (OER) that exceeds 30.0% (this is becoming a *conservative* number as new issues are identified). You say these questions are “based on” prior actual exams—i.e., rather than directly sampled—and you imply that revisions and editing *created* these errors. With all *sincere* and due respect, does this pass anybody’s smell test? It seems highly unlikely that editing would systematically create new errors; that somehow adaptation could degrade quality questions into flawed questions. Specifically, it is unclear how editing and revision could introduce several of the *conceptual* confusions. For example, how could portfolio VaR questions with robust assumptions *abruptly shift* into inconsistent assumptions that misunderstand the definition of beta? How would an adaptation *suddenly unlearn* fundamental features of the cost of carry model? I do appreciate your motivations to resist the suggestion, but it’s a virtual certainty that the actual exams contained these conceptual errors.
- **FRM Exam Development:** Thank you for the high-level description. This appears to be a comprehensive process. From my perspective as a teacher on the “front-line” with FRM candidates (who are on the other side of the exam, so to speak), this is what I perceive to be the key challenge: *aligning this effort with the expectations and study efforts of your customers*. The syllabus (readings and learning objectives) directly informs the efforts of candidates (and EPPs). It’s important that the experts and external reviewers are **not too detached** from this roadmap. Candidates and EPPs want *prospective clarity*: a reliable (to some extent consistent) expectation of what is likely to be tested, and as importantly, what is not likely to be tested (the current syllabus seems ripe for opportunities to squander their precious time). The success of this can be judged by an *extremely common* question we receive: “Will the exam test [something specific]?” behind which is the implicit question “How hard should I study this?” Too often we do not know because there’s no reliable example or guidance, and in many cases, we don’t know because we suspect GARP is not expert in the details.

- My **mere hypothesis about the weakness** in your comprehensive approach (that has obvious and corresponding *strengths* to be sure) is that some or several of the experts involved may be, to various degrees, understandably detached from the syllabus that is assigned to current FRM candidates. There's no inherent reason, for example, to expect an FRM Committee member to be highly versed in Jon Gregory's third edition! (Yet that is what the candidate is studying ... it can't be forgotten).
- **FRM Readings and Curriculum Development:** Your process here sounds good to me. I maintain my suggestions with respect to the desirable properties of Recency, Coherence, and Relevance.
- **Exam Prep Provider (EPP) outreach:** The main thing I want from GARP is what we give to our customers: *respect for our (their) time*. It is extremely time-consuming to prepare materials. GARP has struggled to edit/revise a single robust Practice Exam over a four-year period. By comparison, with a two-person team, we write many more than 180 questions *every year*, in addition to study notes, videos and daily forum support. Our error correction contributions to you have been volunteered (and do not get attributed). It is time-consuming to help candidates clarify their confusion induced by your errors.
- **Accountability and transparency:** Given the high cost and time requirements of the FRM, I would suggest that you plan for increasing levels of accountability and transparency. Many of the details in my memo fall under this umbrella. Your memo describes a wonderfully extensive internal and external effort. But as an EPP who has daily contact with FRM candidates "on this side" of the FRM exam, most of it is simply not transparent to us. For example, no actual, verified exam question(s) have been disclosed to candidates; e.g., it seems reasonable to share selected, actual questions from three or five years ago. One salient manifestation of this current lack of transparency is that we simply do not have good answers to certain questions from our customers; e.g., how will GARP approach this? Will GARP test this? Should I spend a lot of time here, or maybe not much at all?

I hope that's helpful. As long as we offer the FRM preparation product, we will continue to constructively support the FRM candidates who are our customers (and those non-customers who simply visit our forum).

Sincerely,  
 David Harper, [bionicturtle.com](http://bionicturtle.com)  
 5/16/2018