ICEBERG RISK

An Adventure in Portfolio Theory

Final, Final Instalment

Kent Osband
kosband@nyc.rr.com

August 2002

PRE-PUBLICATION REVIEW MANUSCRIPT
EXCLUSIVE TO READERS OF
WILMOTT.COM

TO BE PUBLISHED BY TEXERE IN
AUGUST 2002

TABLE OF CONTENTS

20. A QUESTION OF REBALANCE

It took Devlin two hours to shepherd Regretta through the math. He was unusually clear, for Devlin anyway, and had broken down everything into bite-sized nuggets. Regretta for her part listened very carefully, chastened by her earlier misinterpretation. Occasionally she replicated some calculations herself just to make sure she understood.

“Well, what do you think?” asked Devlin when they had finished.

“I think it’s very impressive.”

“But will Conway think so?”

“I think he’ll be impressed too.”

“That’s a relief. I was worried you two wouldn’t like it.”

“I didn’t say we—I mean he—will like it. Being impressed and liking something are two different things.”

“You think there are too many formulas, don’t you?”

“Too many, and too messy. Practical finance types just don’t get turned on by Taylor series.” *Unlike you*.

Devlin avoided Regretta’s gaze. “That’s what I feared. But what can I do about it, apart from burying the calculations in spreadsheets?”

“Search for ways to simplify.”

“I’ve done what I can.”

“Are you sure? How about trying to reformulate the problem in different ways?”

“Like what?”

“I don’t know. I’ll try to think about it too.”

“Thanks, Regretta.”

“You’re welcome. Now how about your taking a lady to lunch?”

“Is it time for lunch already?” Devlin looked at his watch. “Gee whiz. Lunch is almost over. I better run or I’ll miss it. See you later, Regretta.” He ran off, leaving a bewildered Regretta behind.

…

That night, Devlin slept fitfully. He dreamt he was at one of the church socials his mother would take him to as a child. All the ladies of the church joined in to prepare a feast. Devlin went up with a plate to be served. But the woman wouldn’t serve him. It was Regretta, and she was whispering something to his mother. His mother shook her head. “Sometimes I just don’t understand that boy,” she said. “Aren’t you hungry, son?”

When Devlin awoke, the first thing he thought was that he should have asked Regretta to lunch. Too bad she confused him yesterday by mentioning church ladies. Very odd. None of the church ladies he knew could hold a candle to Regretta. She was prettier than the rest of them put together and twice as fiery. Too bad all Regretta and he had in common was math.

Or was it? Devlin had never really tried talking with Regretta about anything else. Math and finance theory. If that's all she saw in him, whose fault was that? He better take some initiative.

It took Devlin the better part of a day to figure out what he should say and how. He spent most of the next day summoning up the courage to say it. He just couldn't bear the thought of rejection.

Thursday they met again at group therapy. Afterwards Devlin asked if she had time to get together again and to his relief she said yes. As they walked back to the oaks they had sat in before, Devlin cleared his throat and began to speak. "Listen, Regretta, I'm sorry about not asking you to lunch on Monday. I didn't think…"

"That I'm a lady?" Regretta cut in.

"No. I mean, yes, you're a great lady. Just not a church lady. I got confused."

You and me both. "That's OK, Devlin. I think you just get too caught up sometimes in one way of thinking."

"Yes, you're probably right. All that time focused on complicated Taylor series expansions. I should have aimed for a better balance."

"You see that now, do you? I'm impressed, Devlin. What made you realize that was the problem?"

"Nothing in particular. I just started to put things in perspective, and realized that the best way forward is to make amends."

"And keep making amends."

Devlin nodded sadly. "You're right. I guess I'll never stay completely on track. But what I wanted to ..."

"No one stays completely on track, Devlin. Rebalancing is a continual struggle. It's hard to figure out how much is enough, how much is too much. I ran a little spreadsheet for insight."

"Spreadsheet? On what? On the chances of…"

"On some simple rebalancing problems. Here, let me show you." Regretta pulled a chart out of her bag and handed it to Devlin. "Suppose your target weight is 5% while the asset falls 80%. Without rebalancing the portfolio loss is 4%. If you rebalance once in the middle—by which I mean halfway down in log terms rather than a 40% loss—the total loss is 5.45%. With sufficiently frequent rebalancing the portfolio loss exceeds 7.7%."

REBAL ANCING WITH ONE RISKY ASSET: w=5% AND x= 80%

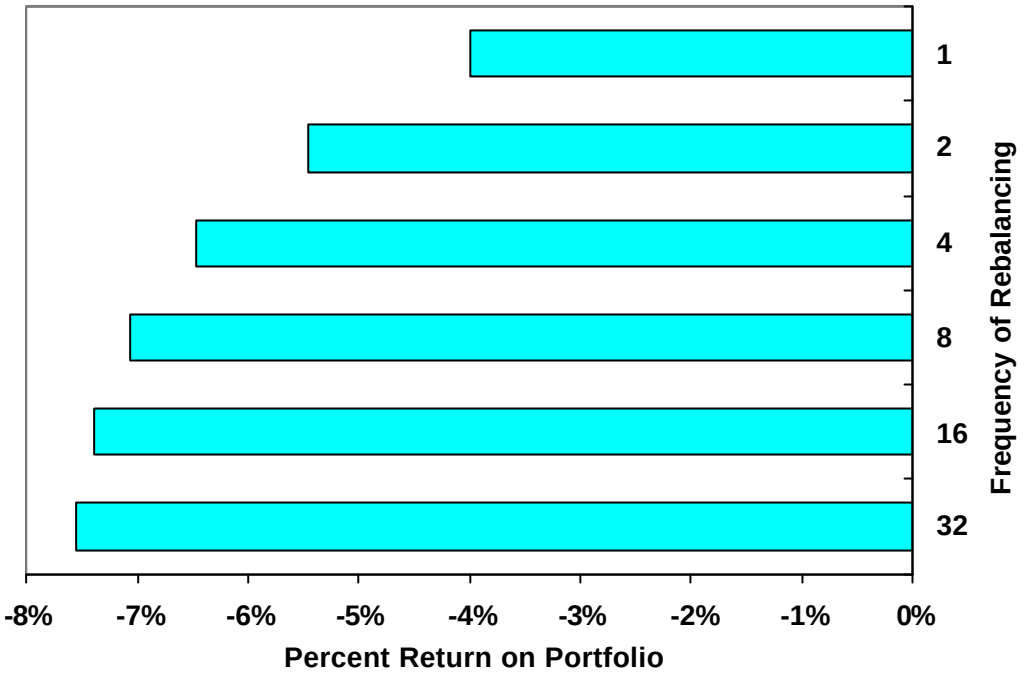

Despite the other things on his mind, Devlin was hooked. “What’s your point?”

“My point is that continual costless rebalancing makes the log of portfolio returns linear in both the log returns and the target portfolio weights. That simplifies your recipes a lot. The conditional log risk-adjusted return simplifies to the log portfolio return less ½c times its variance. And that’s exact rather than an approximation. That’s amazing, isn’t it?”

“It sure is.”

“I tell you, I was really looking forward to explaining it to you. I didn’t realize that you had figured it out on your own.”

“Well, I…”

“Then again, maybe there are a few things I figured out that you haven’t. For example, I thought I ran into a paradox when I let the reset period vary.”

“Reset period?”

“I mean the time interval of observation. To keep things tractable I allow the regimes and portfolio targets to change only at specified resets.”

“And what was the paradox?”

“That everything collapsed to a single-regime case as the reset shrank to zero even though you’re allowing the maximum switching of regimes. But now it all makes sense.”

“Why?”

“Aha! So I did figure out something you didn’t. You see, when regime-switching gets sufficiently frequent, a kind of law of large numbers effect kicks in. Assets start behaving as if they spent all their time in an average composite regime. To make quick resets interesting I

had to introduce Poisson jumps, where there's a constant instantaneous probability of decay. That yielded another neat recipe, which is generally quite easy to calculate. In fact, as long as the Poisson risks aren't too severe, the optimal mix looks like the solution to a Sharpe ratio maximization problem. Here, let me show them to you:" Regretta pulled another sheet out of her bag:

Simplified Recipes

1) With lognormal returns and continual, costless rebalancing, conditional log risk-adjusted returns $\mathbb{C}E_k$ equal $w'\ln(\mathbf{1}+\mathrm{M}_k)-\frac{1}{2}cw'\tilde{\Sigma}_k w$, or equivalently $w'\tilde{\mathrm{M}}_k+\frac{1}{2}w'\left[\tilde{s}_{11k}\ldots\tilde{s}_{nnk}\right]'-\frac{1}{2}cw'\tilde{\Sigma}_k w$.

2) If resets are frequent, then the aggregate $\mathbb{C}E$ approaches $w'\tilde{\mathrm{M}}+\frac{1}{2}w'\left[\tilde{s}_{11}\ldots\tilde{s}_{nn}\right]'-\frac{1}{2}cw'\tilde{\Sigma}w$, where $\tilde{\mathrm{M}}=\sum_k p_k\tilde{\mathrm{M}}_k$ and $\tilde{\Sigma}=\sum_k p_k\tilde{\Sigma}_k$.

3) If the base regime 0 is lognormal Brownian and Poisson vector losses of $100L\%$ occur with frequency ℓ, $\mathbb{C}E$ equals $w'\mathrm{M}_0-\frac{1}{2}cw'\tilde{\Sigma}_0 w+\ell\frac{(1-w'L)^{1-c}-1}{1-c}$.

4) If Poisson risks are moderate, the optimal portfolio mix w^* can be estimated as $w^*\cong\frac{1}{c}\left(\tilde{\Sigma}+\ell LL'\right)^{-1}\left(\mathrm{M}_0-\ell L\right)$.

Devlin looked over the recipes. "That's neat. But in the last recipe, how do we know whether Poisson risks are moderate?"

"It depends on the risk of a jump, the square of the risk aversion, and the cube of the percent of the fund at risk. The sensitivity is a good reason for trimming concentration risks and betas in your portfolio."

"I have a sense that most fund managers strive for that, over and above what standard portfolio theory suggests."

"Yes, I think so too. Our approach—I mean your approach—seems to accord better with practical intuition. But it allows a more quantitative assessment. Here, for example, are some indications of how much log risk—adjusted returns get distorted when we ignore third—and higher-order effects in Poisson jumps, assuming a CRR of 3.

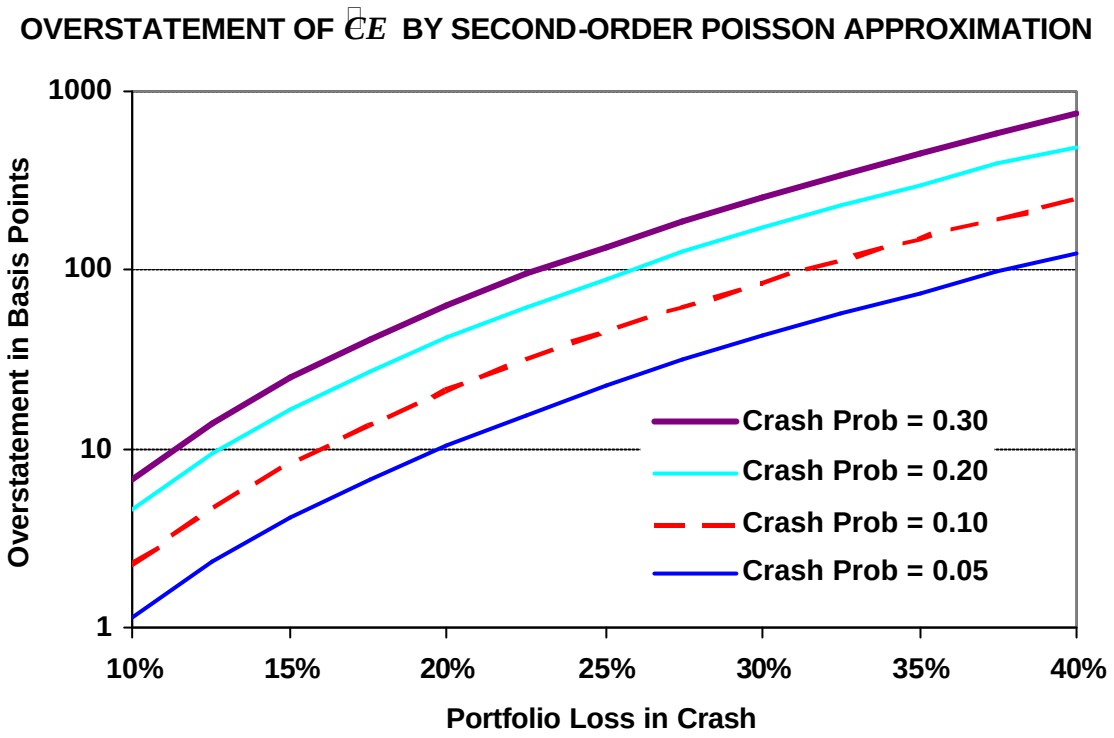

"It's our approach now. But this time you need to explain it more to me. How about your taking a fellow to lunch?"

"What fellow?" asked Regretta, and then smiled. "Yes, please. I'd be delighted."

Now it's Devlin's turn to catch up on the math. It serves him right for foisting all those messy calculations on everyone. I'll try to keep this section a bit simpler than the last.

{NOT POSTED}

21. WEIGHING THE OPTIONS

Devlin would have enjoyed lunch with Regretta even if Club Mad hadn't been serving burritos. She was so lively, so sharp, so insightful. And she looked so good in black. But afterwards something about their conversation began to gnaw at him. "Keep things simple," she kept admonishing him. At the time he thought she just wanted him to trim his explanations. But now he realized she might have been hinting at something deeper. *Was she trying to keep her distance? Trying to let me down gently? Did I come on too strong, too intense?*

Devlin sighed. *That's it. Too intense. Everybody thinks that. Everybody but me, that is. Why do they all get so uptight? If a problem waylays you and keeps you hostage until you solve it, that's not your fault. It's the problem's fault. Granted, I do seem to get waylaid more than most people. But that's just bad luck.*

Bad luck? Get real. Everybody has options. Life isn't some linear trajectory. It's full of options, and you can switch from one branch to another. Maybe it's time to think more about options...

My options. Regretta's options. Our options. There's so much to sort out. Devlin scribbled down some thoughts, twirled them around, and scribbled down some more. At first they didn't amount to anything. Page after page he wadded up and threw in the waste basket.

Remember what Regretta said. Yes, keep things simple. But how? And then it dawned on him. Not exactly simple. But as simple as it could be. He couldn't wait to share his thoughts with Regretta.

Now hold on a second. Why rush over there straight away? To appear even more intense? That might ruin everything. No, I mustn't do that. Not again. Not so quickly. So Devlin kept his thoughts to himself. He kept his distance from Regretta, and confided only in his notebooks.

Devlin's self-control lasted nearly a week. But not quite. One morning Regretta found an envelope under her door, opened it, and began to read:

Dear Regretta,

I'm sorry. I've been trying not to bug you about this, but I can't hold back any longer. We've come a way long way I agree, and I'm thankful for that. But keeping things the way they are won't give us nearly the satisfaction we can get from advancing to the next higher step. So I'd like you to think about various options while I make some proposals.

No doubt you already realize the problem. Options are too nonlinear to fit into conventional portfolio analysis. The best you can do is model them as a fixed fraction "delta" of regular assets, which kind of misses the point. What straight line best approximates a put or call?

Pricing Curves for Ordinary Options

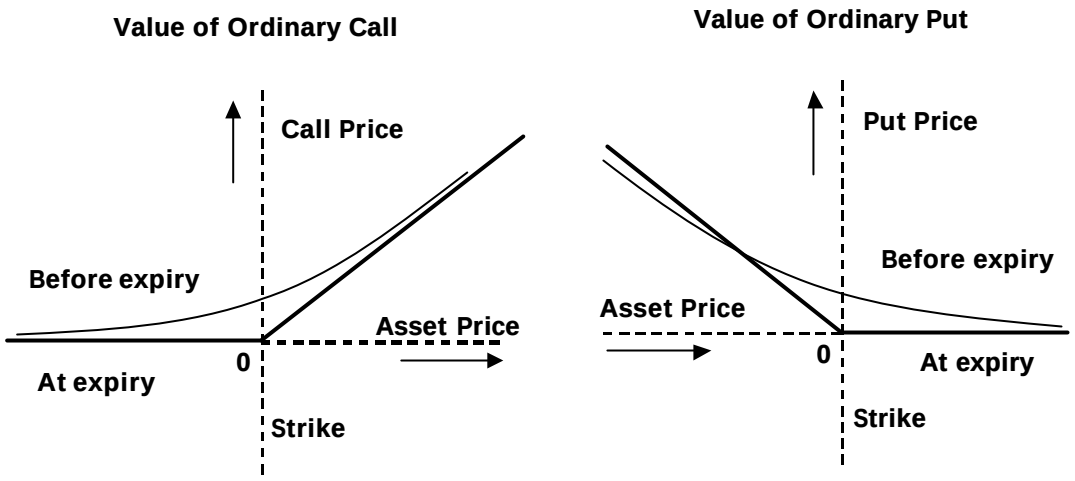

But our framework can handle them. It can handle them in two ways. The first way is by allowing the delta to vary across regimes. A put might have a delta of nearly zero in a boom and nearly one in a bust. The second way is by allowing delta to change even within a regime, assuming a constant rate of change "gamma" that itself might vary by regime.

All you need for calculation, over and above what we already have, are the mean option value and its delta and gamma at the mean of each regime. The general formulas for portfolio risk remain the same; the deltas just change the effective variance. I can't imagine any simpler way to incorporate nonlinearity. And it sure seems to beat the next best alternative because I can't recall seeing one.

Granted, it's not perfect. Gamma's not constant, not even within a regime. Indeed, it can spike very sharply around the strike price close to expiry. To deal with that you should take discrete approximations to delta and gamma, based on average values over a range. This isn't as messy as it sounds: option values at the mean and one standard deviation above and below it will generally suffice for the calculation.

If you can continually and costlessly rebalance, the errors in approximation essentially vanish. One way to see this is to compare our framework's advice with the advice Black-Scholes would give. They're identical, at least in the range where the two approaches overlap. Indeed, you can use our framework to re-generate Black-Scholes. But our approach is theoretically richer, because we allow for multiple regimes and mispricing.

Most likely it's richer in practice too. I won't say "for certain" because even if you can make profits trading options, simple longs and shorts might boost your risk-adjusted portfolio return just as much. Still, I do think it will occasionally help, either by identifying a genuinely great opportunity or by debunking some bogus claim. And it's near essential for a bank or hedge fund that actually uses a lot of options, as opposed to mutual funds that just dabble in them.

Ultimately, what we're really looking for from options is iceberg insurance. To illustrate the difference it can make, imagine your portfolio is perfectly lognormal except for

vulnerability to a sudden "Poisson" crash. The crash is Poisson: it occurs at rate ℓ per unit time and causes a 100L% loss when it does. The chart below shows how much risk premium—that is, fees over and above the expected payout—an investor with a CRR of 3 would be willing to pay for perfect insurance from the crash.

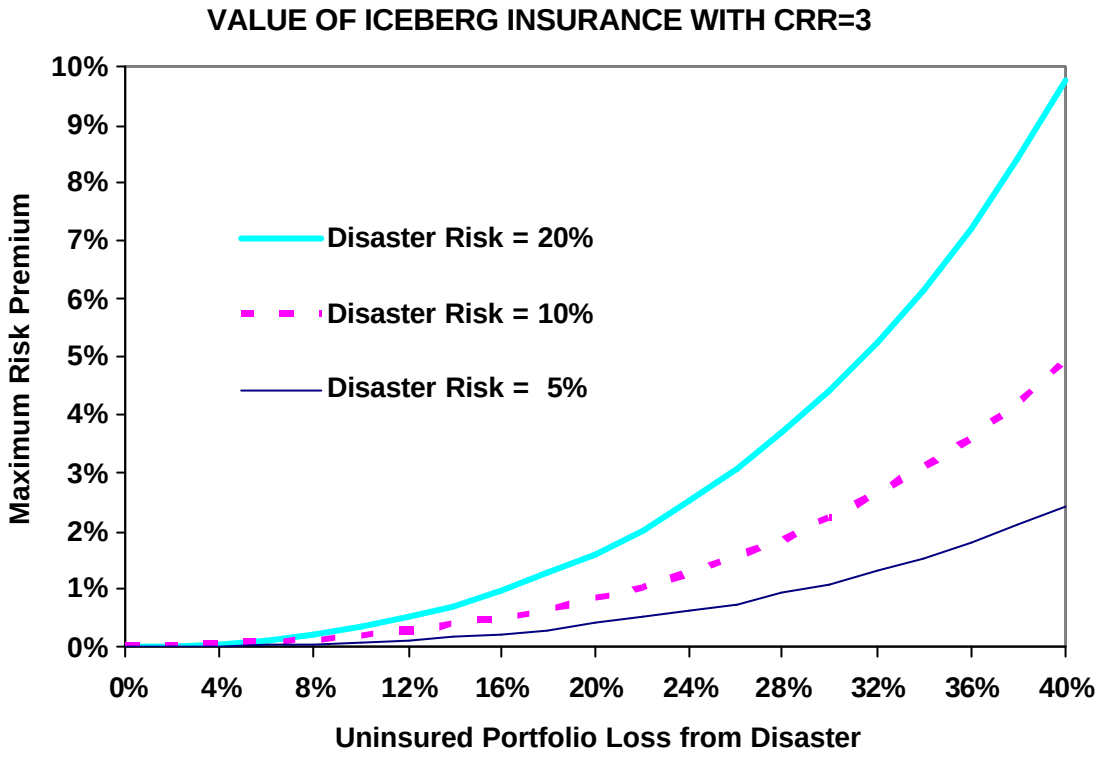

Incorporating options sure looks appealing, doesn't it? But there's one catch. The recipes get messy. Not so messy you can't work with them. Not so messy you can't figure out what they mean: basically just to adjust the portfolio means and variances for the deltas and gammas. But messy enough that some people will just throw up their hands and walk away.

I know you want me to simplify. I tried. I really did. But once you distinguish between assets and underlying price drivers and allow for nonlinear relations between them, there's no way to avoid extra terms, even when you compress things like I did by summarizing all the deltas and gammas in two matrices. Actually, to store the gammas I needed a 3-dimensional matrix. I hope that doesn't bother you.

Anyway, here are the recipes:

Incorporating Options

Let *Y* denote an *N*-vector of nonlinear assets tied to conditionally multivariate normal drivers *X*. Conditional on *X*=M, denote the value of *Y* by $\hat{\Pi}$, the partial derivatives of *Y* with respect to *X* by the *N*×*n* matrix $\hat{\Delta}$, and the second partial derivatives of *Y* with respect to *X* by the 3-D *N*×*n*×*n* matrix $\hat{\Gamma}$, where $w'\hat{\Gamma} \equiv \sum_{h=1}^{n} w_h \hat{\Gamma}_h$. A tilde marks the corresponding values when $\tilde{X} = \tilde{\mathrm{M}}$, while a bar marks the limiting values with continuous rebalancing. Then previous refinements apply subject to the following changes:

Refinement #1 and #2: Recalculate portfolio moments as

$$m_k = w'\hat{\Pi}_k + \tfrac{1}{2}\mathrm{tr}\left(\left(w'\hat{\Gamma}_k\right)\Sigma_k\right)$$

$$v_k = w'\hat{\Delta}_k'\Sigma_k\hat{\Delta}_k w + \tfrac{1}{2}\mathrm{tr}\left(\left(\left(w'\hat{\Gamma}_k\right)\Sigma_k\right)^2\right)$$

Refinement #3: Recalculate $\tilde{A}$, $\tilde{B}$ and Ψ as:

$$\tilde{A}_k = \ln\left(1 + w'\left(\exp\left(\tilde{\Pi}_k\right) - \mathbf{1}\right)\right)$$

$$\tilde{B}_k = \mathrm{diag}\left(\exp\left(\tilde{\mathrm{M}}_k - \tilde{A}_k\mathbf{1}\right)\right)\tilde{\Delta}_k'w$$

$$\tilde{C}_k = \mathrm{diag}\left(\tilde{B}_k\right) - \tilde{B}_k\tilde{B}_k' + \mathrm{diag}\left(\exp\left(\tilde{\mathrm{M}}_k - \tilde{A}_k\mathbf{1}\right)\right)\left(\tilde{\Gamma}_k w\right)\mathrm{diag}\left(\exp\left(\tilde{\mathrm{M}}_k\right)\right)$$

$$\Psi_k \equiv \mathrm{I} + (c-1)\tilde{C}_k\tilde{\Sigma}_k$$

Continuous Rebalancing: Recalculate $\bar{E}$ as:

$$\bar{E} = w'\bar{\Pi} + \tfrac{1}{2}\mathrm{tr}\left(\left(\mathrm{diag}\left(\bar{\Delta}'w\right) + w'\bar{\Gamma}\right)\tilde{\Sigma}\right) - \tfrac{1}{2}cw'\bar{\Delta}\tilde{\Sigma}\bar{\Delta}'w \quad .$$

If you want me to show you the calculations, I'll be glad to. Any time.

Your friend,

Devlin

Any time? Too bad Devlin didn't offer the same services to everyone. While Regretta rereads the letter for deeper meaning, let's try to reproduce the recipes.

{NOT POSTED}

22. FIXING THE FOCUS

The next Visitors Day Conway arrived bright and early. Devlin and Regretta took turns briefing him on the higher-order refinements, on the simplifications that continuous rebalancing can provide, and on the incorporation of options.

"I'm amazed how much you two have accomplished in two weeks," said Conway. *Almost as amazing as how well you two seem to be getting along.* "But I hope you're not expecting me to remember all this."

"Of course not," said Regretta. "Just remember the really important stuff."

"Which is?"

"That the approach is comprehensive," said Devlin.

"But flexible," said Regretta,

"And as detailed as you want," said Devlin.

"Or neat and parsimonious. So you can focus on what you want and pull out the recipe that suits you best."

"Or better yet, have you two pull it out for me" interjected Conway before they both got carried away. "How about I try to bring you two onboard with me at Megabucks?"

"Not me," said Devlin. "I don't want to rejoin a club that wouldn't have me for a member."

"Me neither," said Regretta. "Too many bad memories of life in the fast lane."

"Trust me, the asset management side isn't fast. Half-fast is more like it." Conway smiled, but couldn't elicit one in return. "Look, suit yourself. But I'm going to need some help handling all these assets and regimes. How do you suppose I get it?"

"Simplify," said Devlin and Regretta nearly in unison. They laughed. "Tell him your idea," said Devlin to Regretta.

"You might call it a CAPM-type approach," said Regretta.

"The Capital Asset Pricing Model? What's that have to do with regime-switching?"

"In its original form, nothing. And I'm not talking about market pricing in equilibrium. I just mean the idea of decomposing all assets into common factors and independent residuals."

"What's the point of that?"

"If you can do it, it saves tons on information needs."

"How so? Don't you need even more parameters to track all the alphas and betas?"

"Sure, but then most of the covariances vanish. They're all zero except for the covariances between the market indices themselves. And all the alphas and betas and residual variances are assumed not to vary across regimes. Here's a chart illustrating how much fewer parameters are needed, not counting what you'd need to handle options. The contrast was so great I had to put it in log terms."

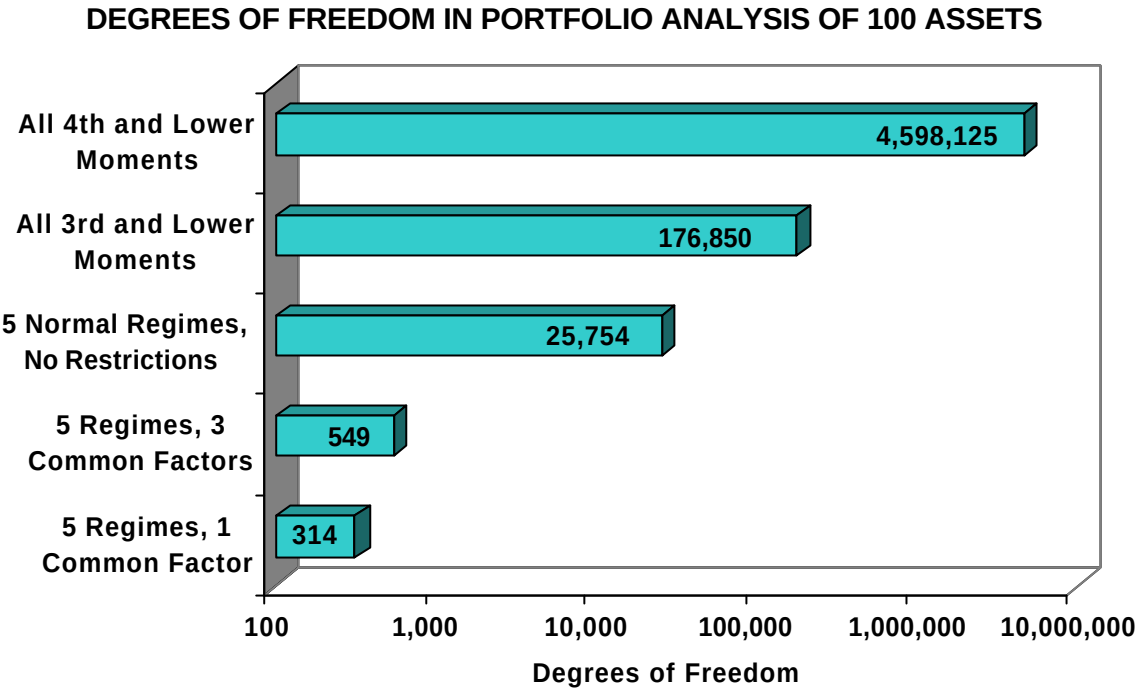

“Those really are huge savings. But doesn’t that eliminate iceberg risk?”

“Hardly. What you’re really worried about is a huge chunk of your portfolio crashing together. Barring some extraordinary concentration of assets, that won’t happen unless one or more core indices crash too. So you may as well focus your iceberg watch on those core indices. Let their means vary a lot across regimes, and covariance matrices too, and most important try to get the regime odds right.”

“What happens to all the other assets?”

“Well, the parts of the assets that track core indices get included with the core indices. As for the residuals, they get treated like independent normal variables. You try to maximize their Sharpe ratio using the standard rules, regardless of what happens in the rest of the portfolio. Convenient, isn’t it?”

“Definitely. But I’ve seen convenience cause trouble before. Devlin, what do you think? Don’t the errors that are bound to creep into the approximations bother you?”

“Yes, they do bother me. But multiplying the number of estimated parameters fifty-fold won’t get rid of them. I’d rather do sensitivity analysis on a few key parameters, or better yet model the uncertainty about them directly.”

“You mean like Black-Litterman does?”

“Not exactly. I don’t dilute views with the market consensus because when there are multiple regimes it’s hard to impute what the consensus is. Instead I model doubts as extra noise-making random variables. Sometimes you can calculate their impact exactly; sometimes you just approximate it.”

“And are five regimes enough?”

“Maybe not enough. Maybe more than enough. It depends on your concentration risks.”

“You mean clusters in one sector or market?”

“Partly. More generally I mean any common vulnerability. If Korean stocks that you own start moving in tandem with Nasdaq, you’ve got extra Nasdaq concentration risk regardless of the sector or market listing.”

“Isn’t that what the multiple regime specs are supposed to capture?”

“Supposed to, yes. But it’s easy to miss the iceberg for the ice cubes. I’m beginning to doubt you should model more than a handful: say, BULL and BEAR or WALK, FLY and DIVE.”

“Then how would you capture the multiple vulnerabilities of Korean stocks to Korea-specific shocks, to dollar/yen, to Nasdaq, and to the global market as a whole?”

“Multiple common factors can handle that without necessarily requiring a separate regime for each blow-up risk. If your absolute portfolio beta with respect to Korean risk is less than 0.05, for example, I wouldn’t generally bother to model it separately.”

Conway shook his head with mock disapproval. “You’ve let me down, Devlin, old boy. I was expecting higher standards. Next thing I know you’ll be advocating absolute concentration caps on everything, just like an old-fashioned risk manager.”

“Concentration caps are primitive. And I don’t like the word ‘absolute’. But a lot of what the new theory advises looks like a combination of Sharpe ratio maximization and concentration caps.”

“Which is kind of what happens in practice.”

Devlin shrugged his shoulders. “Are you sure?”

Touché, Devlin. “No, I’m not sure. Maybe they’re just trying not to fall too far below the benchmark, whatever that happens to be. I guess your theory doesn’t cover that.”

“Not exactly,” said Regretta. “But there’s a quick fix that I think will do the trick.”

“Really? What is it?”

“Just remeasure returns relative to the benchmark rather than to the risk-free rate.”

“Say, that is easy. But how does the new optimum compare to the old?”

“Think about it. You’ll be looking for a portfolio that’s optimal after you subtract off one unit of excess returns on the benchmark. So that means you just take the previously optimal portfolio and add to it one unit of the benchmark or its proxy funded with borrowed T-bills.”

“Thanks, Regretta. But what if the portfolio managers don’t agree?”

“Then maybe you should score the portfolio on its performance.”

“Don’t you want to take risk into account?”

“By performance I mean the utility of returns, not the returns themselves. In the long run this should average out to expected utility if forecasts are correct.”

“I’m afraid utility won’t be very intuitively appealing. Not everyone will accept the framework.”

“Then convert average utility to a certainty equivalent. That’s just the risk-adjusted return. Everybody in the business claims to understand that.”

“I’m not sure they do”.

“They will if you start factoring it in to their rewards.”

Conway nodded. “Fair point. But these utility measures apply to the portfolio as a whole. How am I supposed to score individual performance in a multi-manager fund?”

“There’s no perfect way to do this. But on the whole I favor measuring marginal contributions. Something like the risk-adjusted return of the whole fund less what the risk-adjusted return would have been if the manager hadn’t been there.”

“If the manager hadn’t been there, a lot of things might have changed. There could be more than one measure of what might have been.”

“Which will mean more than one measure of managerial performance. So be it. But overall you want to encourage managers to add diversified alpha in normal times without aggravating iceberg risk in crises. That’s what measures of marginal contributions try to do. Here’s an example I charted for a bivariate normal world in which the rest of the portfolio has a Sharpe ratio of 0.5. If the correlations didn’t matter the lines of equal marginal contribution would be horizontal. Instead, for the most part they’re very steep.”

MARGINAL CONTRIBUTIONS TO RISK-ADJUSTED RETURN (*CE*)
(Baseline Sharpe=0.5, CRR=3)

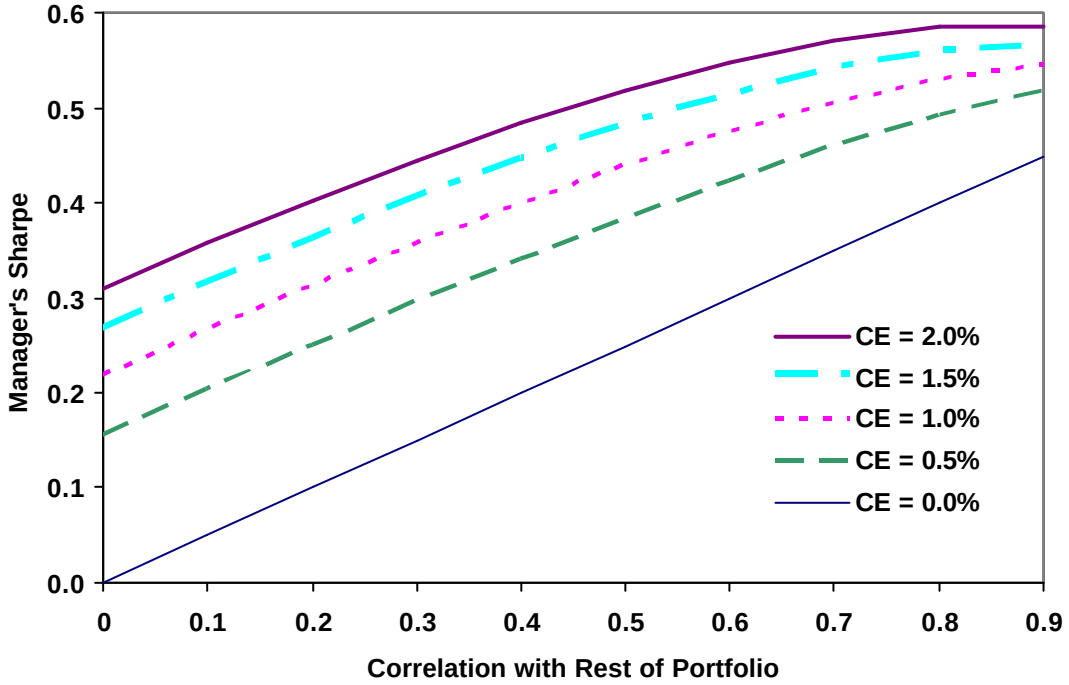

“Yes, I understand,” said Conway. “I was using Sharpe ratios to make the same point. How do your measures compare with that?”

“In a multivariate normal world the certainty equivalent will be proportional to the aggregate Sharpe ratio squared. But it will penalize more for iceberg risk than the Sharpe squared will.”

“Great. This is all very clear, Regretta. Thanks.”

Devlin suddenly spoke up. “Well it may all be very clear to you two but it’s not to me.”

“What’s the matter?” asked Regretta. “I thought you agreed with performance scoring. We talked about this earlier.”

“I did agree then. And in principle I still do. But while you were explaining things to Conway I got to thinking about reliability.”

“Reliability of what?”

“Of the average utility scores. How much are they likely to deviate from the expected utility?”

Regretta reflected a moment. “I suppose it depends on how noisy the portfolio is–volatility and all that–and on how long you evaluate the portfolio.”

“So why don’t we try to estimate the reliability and factor that into the evaluation?”

“I don’t get it,” said Conway.

“I do,” said Regretta. “The actual performance scores might be thought of as random observations on the true expected performance score. I’ve been arguing that in the long run the average will converge to the expected value.”

“Which it should,” said Conway. “Do you disagree with that, Devlin?”

“No I don’t disagree. In the long run Regretta’s method should work fine. But in the short run we need to be wary. You wouldn’t fire a manager or carry out a major portfolio reshuffle solely on account of one bad trading day, would you?”

“Not unless the day were really extreme. Even then I’d be tempted to dismiss it as an outlier.”

“Agreed. But the long run is only a chain of daily results. So each result carries some information value and I’d like to know how much.”

“Ah, I see. That is a good question. But if you don’t mind, let’s defer it to another time. My brain is overloaded. It can’t handle any more.”

“How about I sign you up for a session at the spa?” asked Regretta. “It’s one of Club Mad’s many comforts.”

“That sounds wonderful. You know, I’m starting to get envious of you two.”

While Conway takes on some steam, let’s probe some more into CAPM-type reductions and performance scoring. But not too deeply: Conway’s isn’t the only brain getting overloaded.

{NOT POSTED}

23. THE ROOSTER PRINCIPLE

For weeks Conway had kept mum about his outings to Club Mad. He had been far from sure they'd be productive and didn't want to be judged by the company he kept. But now he had something to crow about. Devlin and Regretta had delivered a portfolio analysis scheme far beyond his expectations.

Surely Jim will appreciate this. But when Conway sketched the new approach, Jim was cautious. "It's intriguing. Very creative and possibly useful. But I'm afraid a lot of it is over my head."

"A lot of it is over my head too. Fortunately I had friends to help me with the math. We can program in the calculations so they're done for us."

"I don't feel comfortable using things I don't fully understand."

"You drive a car and use computers. Do you fully understand internal combustion engines and silicon chip semiconductors?"

"No, but I understand what I need to do to use them, because I watched a lot of other people use them first. Here you're asking me to be a pioneer. I don't think I'm up to it."

"Somebody has to start. Why not us?"

Jim leaned back in his chair and grinned. "One thing I've learned in finance is the importance of being second. The pioneer is the guy with the arrow in his back."

Conway laughed. "Well put. Only this is different. It's designed to reduce our risks, not raise them."

"Even if I believe you, others might not. Remember, we have to manage the appearances of risk, and not just risk itself."

"Yes, your brother explained that to me once. I haven't forgotten. I just want a chance to win people over. Some experiment, some forum, some…" Conway's voice trailed off. "I don't know," he said quietly and lowered his eyes.

"Conway, let me give you some friendly advice," said Jim in a fatherly tone. "Never hold strong convictions in finance. They'll just be a burden to you."

"I never suffered them before. But this one's got hold of me."

"I can see that. OK, Conway, here's what I'll do for you. I'll invite all the portfolio managers to a meeting on 'Iceberg Risk and How to Deal with It'. I won't make them come and I won't make them stay. But I will give you a forum to make your case. Depending on the reception, we can discuss launching an experiment afterwards."

Conway's eyes brightened. "Thanks, Jim. I really appreciate the opportunity. You won't regret this."

"I hope you won't either. Is Friday afternoon too soon for you?

"No, that will be great."

…….

Conway spent the next few days preparing. The challenge wasn't to cram everything in but to strip everything inessential out. *I have to show them how simple it really is.* Somehow he managed. By Friday morning Conway was quite satisfied with the presentation and confident of success.

Jim opened the meeting with a brief introduction. "Thanks for coming, everybody. It suggests we're all interested in the same question. It also suggests we're still looking for answers. Conway here claims to have some. He shared them with me and I was impressed. Impressed both by the parts I understood and the parts I didn't." Jim smiled and drew a few chuckles. "So I asked him to run thru it again with a smarter audience—namely, you—while I check out whether I understand it better the second time. Now I want you all to lend him your ears, while Conway tries to repay with interest."

Conway walked over to the slide projector. "Thanks. I want to show you all a new way to analyze portfolio risk. It does everything the old way does and more. It's also the simplest possible way to extend the old way without violating theory or common sense. I'll try to be clear, but if I'm not feel free to interrupt with questions."

A hand shot up. "Is this going to involve a lot of math?"

"No. Just a smidgen. The new way is based on a theory that does involve a lot of math. It's more complicated than the standard theory and requires more computation to apply. But I say let's leave the computation to mindless computers, so we can focus on what you need your noggin for. In fact, a good part of what I want to say can be summarized in two pictures. Here's the first:

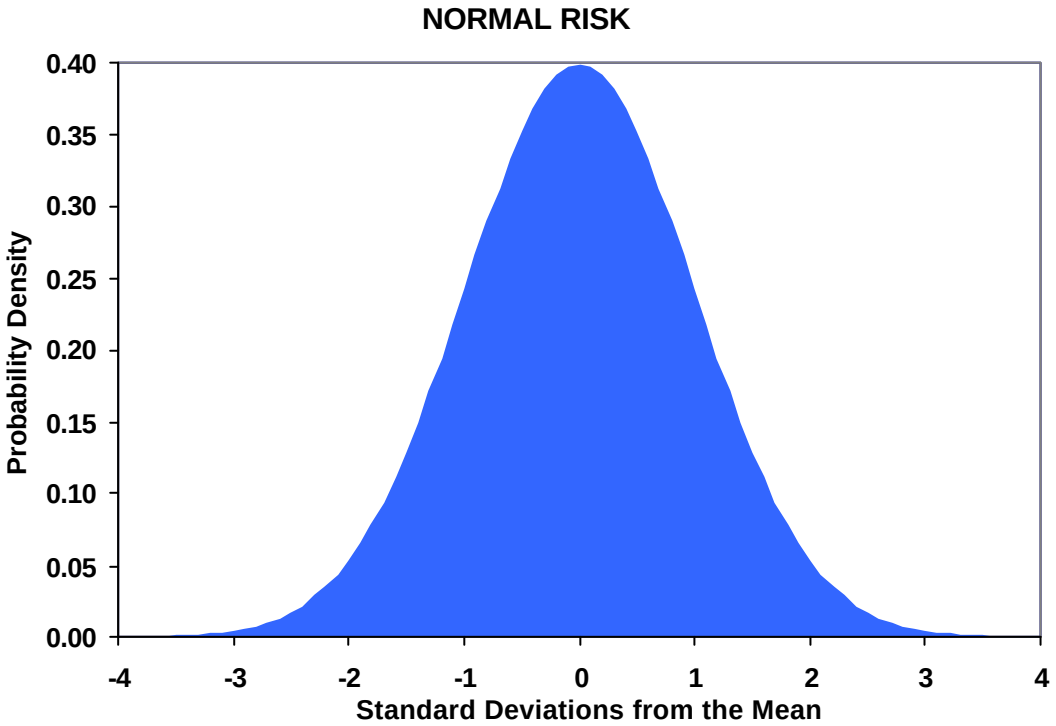

"This is a picture of normal, bell-shaped risk: 95% of the action occurs within 2 standard deviations of the mean; occasionally you go out to 3 standard deviations, and 4 or more standard deviations you can basically forget. Now I know you're all familiar with this. But I'm showing you all again to remind you of something very important: this is the *only* risk that the standard way looks at. Yes. The standard way assumes that every risk in every portfolio

looks like this. The only differences are the values for the mean and standard deviation, which you calculate using standard formulas.

“Now, if you have to choose one risk to look at, this is the right one to choose. Why? There are two reasons. The first is that if you average out a lot of independent variables the risks always look normal, except maybe in the extreme tails. That was proved centuries ago in the Central Limit Theorem, the most famous theorem in statistics. The second reason is that a lot of correlated risks can be viewed as independent if you just change the way you measure them, so that you can apply the Central Limit Theorem to them too.”

“However, often vast blocs of assets share a common driving factor. For example, most US stocks ultimately depend on the health of the US economy, so if the latter falters they’ll falter too. Some will falter more, some less—economic growth may affect some stocks more than others, and in no stock is growth the only influence. Still, that group will share a common vulnerability.

“When asset returns have a common risk factor, they can’t be viewed as independent no matter how you try to slice them or dice them. In that case, a portfolio’s risks won’t be normal unless the common risk is normal.”

Another hand went up. “What if the common risk is approximately normal, or if it’s small relative to other common risks?” asked a woman.

“Good question. Yes, in that case, the standard approach will be good enough. With regard to US economic growth, for example, I’m inclined to treat it as an approximately normal influence. But some driving factors aren’t even approximately normal.”

“Like what?”

“Like the influence of the market as a whole. If Nasdaq crashes, then most of the stocks will crash with it.”

“Of course. If they didn’t, Nasdaq wouldn’t crash. Isn’t that a circular argument?”

“Yes and no. I suppose it would be more precise to say that Nasdaq crashed because some combination of a common shock and contagion effects made a lot of stocks in Nasdaq fall together. So the true common factor is the shock and/or contagion, which I’m proxying by the Nasdaq index itself. Having said that, I don’t see any harm in identifying a common factor with its index proxy. In fact, in some ways it’s better, because we can mitigate a risk by shorting its proxy, even we can’t trade the risk itself. Does that answer your question?”

“Yes.” The woman nodded.

“Good. Now I don’t think I need to convince anyone here that Nasdaq risk isn’t normal. Not after what we’ve seen the last two years. Its spectacular rise and fall might be likened to the launch and sinking of the Titanic. Iceberg risk. Only it’s a lot more prevalent in finance than in shipping. How can we amend standard risk analysis to incorporate iceberg risk?

“That’s a severe challenge. To begin with, what assumptions should we make about the probability distribution of iceberg risk? Does it mean lumpy tails or a general thickening? Is it skewed downward, and if so by how much? And how do we incorporate the likelihood that we’re not even sure exactly what the distribution is? For iceberg risks nearly always come shrouded in uncertainty. In fact, the more I think about it, there’s only one good way to summarize iceberg risk in a picture. Here it is:”

ICEBERG RISK

Conway waited for the laughter to subside and then continued. “How do you make a useful model out of the notion that anything might happen? A friend of mine wrestled with this for months before he found the only tractable answer. Let the common factor or factors determine the general state of the world, also called the ‘regime’. Within each regime assume everything is normal. But the parameters of different regimes and their probabilities of occurrence can vary as you wish. That’s called ‘conditional normality’. Graphically it amounts to an overlay of different bell curves.”

“Any risks you’re interested in can be modelled this way. But the important thing is that a handful of regimes typically suffice to capture your main concerns.”

Henry spoke up. “Conway, why do call this is the only tractable answer? There’s a host of distributions you could choose from. Some of them capture fat tails a lot easier than normal distributions do.”

“Fat tails of individual assets, yes. But in portfolio analysis we’re not nearly as concerned about fat tails of individual assets as we are fat tails of portfolios. There’s not necessary connection between the two. A portfolio of high-yield bonds, each with huge default risk, might look very normal in aggregate if you diversify enough and hedge out market risk. Conversely, assets with virtually no tails might all be vulnerable to the same common risk, causing a huge fat tail in the portfolio.”

“So model the correlation too.”

“Correlation alone can’t capture the odds that a lot of things tank together. Not unless every asset and combination of assets is normal. With any other distribution you need to specify higher-order moments and cross moments. The co-skewnesses and co-kurtoses of an average-size portfolio easily number in the millions. So you may as well model the conditional dependence directly, and conditional normality is the easiest way to do this.”

“I need to think about that one,” said Henry.

“When I first heard it I didn’t believe it either. Now it seems obvious. Come by my office afterwards and I’ll show you the evidence that turned me around. In the meantime, whether you think I could do this other ways or not, do you understand what I mean by conditional normality?”

“Sure,” said Henry. “Each regime is fully normal, but the means, variances, and covariances are liable to change.”

“Exactly. A regime is something like “bull market” or “bear market” or “liquidity squeeze”. It’s what most people regard as the big market picture. Actually, if you strip away the math, that’s the natural way of thinking about big risk. Standard mean-variance analysis focuses more on little risks. Sometimes that’s enough. Sometimes it misses the iceberg for the ice cubes.”

“No approach can completely avoid icebergs,” said Henry. “How do you weigh the risks against the foregone rewards?”

“Another good question. The truth is, conditional normality on its own won’t get you anywhere. You need a scoring system: some way to compare normal risks, iceberg risks, and rewards so that you can judge which portfolio is best.”

Another hand went up. “Do you mean something like Sharpe ratios?”

“Yes. But Sharpe ratios ignore iceberg risk, since they look only at mean and variance. And since Sharpe ratios are independent of leverage, they can’t tell you how much leverage to apply. You need a more sophisticated measure.”

“How can you decide something like leverage without knowing the investor’s tolerance for risk?”

“You can’t. At the same time, we don’t want our scoring rules to demand too much information about investors’ risk tolerance, because we rarely have much information at hand. For example, it would be nice if, all else being equal, the optimal percentage allocations of a portfolio don’t depend on its absolute amount. Let’s also assume our investors aren’t suckers: that they will never knowingly take bets guaranteed to lose them money. Or at least that they don’t want Megabucks to take sucker bets on their behalf. That doesn’t sound too restrictive, does it?”

Conway waited for objections but none came, so he continued. “Great. We all seem to be on the same page. And guess what? Under the conditions I laid out, economists have proved that there’s basically only one kind of scoring rule to use. It’s the expected value of a power function of the investor’s wealth, with a sign chosen to ensure that more is better than less. I’ve written it out for you in this slide:

$$\textbf{Given a coefficient of relative risk aversion } CRR \neq 1,$$

$$\textbf{maximize expected utility } EU = E\left[\frac{Wealth^{1-CRR}}{1-CRR}\right].$$

“Actually you could add any constant, or multiply by any constant, and get the same implied behavior. And when the *CRR* equals one you should replace the power function with a logarithm. Apart from that it’s unique. ‘Expected utility’ is just economists’ name for a scoring rule that people appear to maximize even if most likely they’re not consciously doing so. The $E[\cdot]$ is just the statistical symbol for expectation given the relevant probability distribution on wealth. As for the *CRR*, underneath its long-winded title it’s just an ordinary number: zero if you don’t mind risk, positive if you do.

“There are two tricky parts to applying this formula. The first is deciding the definition of wealth. Any economist worth his salt will tell you that wealth includes not just liquid assets but also real estate and human capital; that is, the discounted future stream of earnings from

employment. For most people human capital is the biggest part. But we finance types love to simplify so let's just restrict wealth to portfolio wealth. In fact, to keep things tidy we'll often just restrict this to the portfolio wealth we happen to manage.

"The second tricky part is to set a CRR appropriate to the definition of wealth. If you assume that most investors more or less correctly apply this approach, then some calculations I won't try to defend here suggest using a CRR in the range of two to four for the stock market as a whole. Alternatively, you can directly ask investors some loaded questions about their willingness to take big bets. If a bet offers even odds of doubling your total portfolio or halving it and you won't take the bet, then your CRR is at least 1. If you would be willing to risk half of your portfolio for a 94% chance of huge gain, your CRR is less than 5.

"It's an interesting parlor game to ask the same person different loaded questions and check how consistent her answers are. More to the point, it would be useful to ask these questions of each other, with the reference being not our personal wealth but how we think our clients would like us to manage their funds. Let's see whether we share roughly the same opinions and check whether our clients agree.

Several people stirred in their seats and a few hands were raised. "That might be premature," said Jim. "I don't want to confuse our clients or get them thinking we've lost our way." Others murmured their approval and the hands went back down.

Uh-oh, they're awfully touchy. "Yes, perhaps we should just keep it among ourselves. Besides, nothing keeps us from experimenting with different CRR values and seeing what difference it makes. So with your permission I'd like to proceed to the next step, namely how to score conditionally normal portfolios. Are you all still with me?"

"Wait a second," said Henry. "Normal returns are unbounded. But your expected utility is defined only for positive wealth. So you can't legitimately mix the two."

Conway wagged his finger at Henry as if to scold him. "Shame on you, Henry. If it weren't for you I would have made a clean getaway. But you're right. To reconcile the two concepts you have to either shift from normality to lognormality or ignore the negative terms. It makes the math a lot hairier and forces you to accept some approximations. Which approximation to choose partly depends on how frequently you rebalance your portfolio to restore target weights."

"Am I the only one who finds this confusing?" asked Jim. Other voices replied. "No." "Me too." "This is over my head."

I'm losing them. "You're right, Jim, and the rest of you too. It is confusing. That's why I want to leave the details to computers. What I want to present here is just the broad structure, with a formulation that captures the essence and ignores the rest. I think you'll find it illuminating."

"Well, I'd like to see it," said Jim. "But there's no point to keeping people who feel they've had their fill. If any of you want to leave, go ahead."

Nearly half the people left. Conway pretended he didn't mind. "I'm not surprised we've thinned out. Let's face it: Expected utility isn't the usual way of thinking about returns. It does have the attraction though that the expected utility of the whole averages out the expected utility of its parts. That is,

If each regime k has probability p_k and conditional expected utility EU_k,

$$EU = \sum_k p_k EU_k \quad .$$

"Next let's convert expected utility into something a bit more natural, namely the guaranteed return that would yield the same satisfaction. Economists call that the 'certainty equivalent'. In other words:

Define the certainty equivalent CE such that $|CE + EU_k| \circ \left(\quad \right)^{-1}$.

"The brackets around the *EU* denote absolute value. I imposed them so as not to worry about the sign. I also divided thru by baseline wealth to convert everything to percentage returns. The reason I did this is that there turns out to be a very simple and intuitively way to estimate *CE*:

$$CE_k \cong \textbf{ConditionalMean}_k - \frac{CRR}{2} \times \textbf{ConditionalVariance}_k \quad .$$

"In other words, the risk-adjusted return in each regime approximately equals the mean in that regime less a multiple ½CRR of the variance in that regime. Note that the penalty on variance rises directly with your risk aversion, which makes sense."

Conway paused for feedback. "Is that it?" asked Henry.

"Basically. You need some small adjustments to compensate for the differences between logs and percentages and expected slippage in portfolio weights. Actually it's better to convert everything to logs. Among other things that lets you reinterpret portfolio optimization as a kind of entropy minimization. It makes finance look more like physics. But I'm trying to keep things simple here."

"Where's the formula for the optimal portfolio mix?" asked Henry.

"There is no explicit formula except for the special case of one regime. Then it reduces to standard Sharpe ratio maximization plus a rule that makes the absolute weights on risky assets inverse to the CRR. The optimal portfolio for the general case also represents a kind of Sharpe ratio maximization, provided you replace the regime probabilities with risk-adjusted probabilities."

"What kind of adjustments are you talking about?"

"The adjustments give more weight to the regimes with lower expected utilities and make you more sensitive to iceberg risk."

"How much more sensitive?"

"That depends on your CRR and the specific risk/reward tradeoffs. You wouldn't want it any other way."

"Can your approach handle options?"

"Yes, it can. You just need to feed it estimated option values, deltas, and gammas at the means of different regimes. In contrast, standard mean/variance analysis has to pretend that the option delta is fixed, ignoring all the nonlinearities that distinguish options from ordinary securities."

“Gee, I don’t know about this,” said Jim. “You said you tried to economize on information. But it looks like overload to me.”

“It doesn’t have to be. You can reduce each asset CAPM-style to a random deviation around a beta-weighted sum of some core factors plus a constant. Assuming these random deviations are independent, you can focus on the correlations and iceberg risks you think are most important, with minimal information needs and distraction.”

“I still doubt most of our managers and analysts would take the information requests seriously. We’re practitioners, not a bunch of theorists.”

“Then maybe we should keep track of how much everyone contributes to risk-adjusted returns and use that as part of their evaluation. Practical folks tend to understand a performance objective pretty well when it’s tied to money.”

“What do you mean by that, Conway?” Jim’s voice took on a harder edge.

Did I say something wrong? “Nothing, really. I’m just pointing out that by scoring the utility of performance and converting it to risk-adjusted returns, you can measure not only the performance of the whole but also the marginal contributions of its parts.”

“Oh really?”

“Yes. For example, you might measure a manager’s marginal contribution as the difference between the risk-adjusted return of the whole portfolio and what it would have been without the manager—say, replacing the manager’s actual subportfolio with some benchmark index. In principle, we could measure everyone’s performance objectively, without regard to rank or seniority.”

“I thought this meeting was supposed to be about iceberg risk,” said a grey-haired man. “What’s that have to do with performance measurement?”

“I would think they ought to have a lot to do with each other,” said Conway. “For example…”

The grey-haired man cut Conway off. “Jim, we seem to be getting highly speculative here. Why bother fixing what isn’t broken? Granted, we did go down with Nasdaq over the past 15 months, but so did a lot of other people; we didn’t lose that much more than the market average. I resent the insinuation that something is wrong with our incentives.”

“Likewise,” said someone else. “If you ask me, the very notion of introducing a new system to analyze iceberg risk involves a lot of iceberg risk. Who knows what new problems could arise?”

“Here, here,” said another.

Before Conway could defend himself, Henry asked a question. “Conway, have you done any calculations of the errors likely to arise from using ex-post performance as a proxy for expected utility?”

“Well, no,” replied Conway, awkwardly, “but I never meant to suggest that one should reward solely on that basis. It would need much more work to …”

Jim cut in, “Yes, I think we can all agree on that. It’s always nice to end a discussion with consensus. Conway, on behalf of everyone here, and I’m sure everyone who left as well, I’d

like to thank you for your very stimulating presentation. It's good every few years or so to get this kind of abstract theoretical perspective."

Others nodded in assent and then left quickly, whispering to each other. Conway and Jim were left alone. Conway buried his face in his hands.

"Disappointed?" asked Jim.

"Shattered." Conway looked up. "What went wrong?"

"You violated the most important principle in practical finance: the Rooster Principle."

"What's that?"

"The Rooster makes all the rules."

"Who's the Rooster?"

"You mean you don't know? That's another violation."

"So it's you."

"Me, some of the senior managers, or any finance hotshot. The folks at the top of the heap. We make the rules. Not some math theory."

"Don't you believe in merit?"

"Sure I believe in merit. But there are many kinds of merit. There's merit in managing big egos and getting them to work well with each other. There's merit in attracting new clients. There's even more merit in convincing the clients you already have that your skill accounted for most of their gains and bad luck for most of their losses."

"What's meritworthy about that?"

"It soothes them. Skill and luck are woven together so tightly in finance. Why should clients torture themselves trying to unravel them? Instead I help them feel better both about their investments and themselves." Jim spoke without a trace of sarcasm.

Conway was puzzled. "I don't get it, Jim. If that's how you feel, why bother hearing out my presentation? Why even hire me?"

"Because the Rooster Principle has a qualifier: One day a rooster, the next day a feather duster."

Conway laughed. "Perhaps this principle is deeper than I thought. So I take it you do want to monitor iceberg risk after all."

"It's not a priority right now. We're still reeling from the Nasdaq iceberg we hit. Still, I learned two very important things from your presentation."

"Namely?"

"First, that looking in more sophisticated ways at financial risk can potentially add a lot of value. Second, that this will never be this group's forte."

"You must find that very discouraging."

“Far from it. One of our holdings is a giant financial risk trader. It hires hundreds of PhDs to churn out analytic models of risks and identify arbitrage opportunities. On that basis it makes huge proprietary bets. I never took a big position before because I didn’t really understand what they’re doing, and you know how I feel about investing in things I don’t understand. But your presentation today changed my mind. That combined with the fact that the stock is trading well under half of last year’s highs.”

“Are you speaking about End Run?”

“Indeed. I’m going to make it our biggest single position: 10% of our total portfolio. I realize that’s over our stated limits. But I figured out some creative accounting that will allow it for a while. Long enough to make up some lost ground.”

Conway gulped. “There are some rumors going around that End Run can’t provide the liquidity to clients it used to, which undercuts many of its proprietary bets, which in turn further constricts the liquidity it can provide.”

“I know that. That’s partly why the stock has dropped so much. But you’ve made me confident that smart guys will work things out.”

“Long-Term Capital had at least as smart guys and blew up anyway.”

“But End Run has Long-Term Capital’s experience to draw on. In fact I wouldn’t be surprised if they’re developing iceberg risk models just like yours. Or better. No offense, Conway, but a hundred heads have got to be better than one.”

“Actually it’s been three heads. And two of them are quite unusual.”

“Conway, stop. One or three, what’s the difference compared with hundreds? Don’t be jealous; I’ll give you full credit for having inspired this idea. Besides, End Run can’t possibly lose us more than 10%.”

“Well, actually, if you keep rebalancing you can lose far more than 10%.”

“You’re obviously much too wound up, Conway. 10% is 10%. The weekend’s coming; why don’t you take off Monday as well to relax? Now, if you’ll excuse me…”

Conway was left alone.

Later that evening Devlin and Regretta were sitting in the Club Mad lounge. “What are you thinking about, Regretta?” asked Devlin.

“I was just wondering how Conway fared in his presentation today.”

“I’m sure it went well. He’s got great communication skills. Not like me.”

“Don’t sell yourself short, Devlin. Without you Conway wouldn’t have had anything to communicate.”

“Without you neither of us would have had anything to communicate. Thanks for teaching us about utility functions.”

“You’re welcome. But the key step was how you combined CRR utility with conditional normality. That was ingenious.”

“That was primitive. And my higher-order follow-ups were so convoluted. I should have gone right away to continuous rebalancing the way you did. Now that was ingenious. Elegant too.”

“And nearly trivial compared with your derivations on options.”

Devin blushed. “I have to admit; I kind of surprised myself on that one. But you know the part I liked best? Working with you on the alpha/beta reductions and the scoring rules.”

Regretta smiled. “I liked that part best too. And Conway sure seemed to appreciate it.”

“Well, I owed him one. The last time I fed him ideas for a presentation, they got him in such hot water he lost is job. I felt terrible about that. That’s why I checked into Club Mad. Well, that and general dismay over my ignorance. Whoever said ‘ignorance is bliss’ didn’t know what he was talking about.”

“I’m happy for you Devlin. Happy and envious.”

“What are you talking about, Regretta?”

“Don’t you see? You’re cured. You’ve made amends to Conway and you’ve overcome your ignorance. There’s no longer any reason for you to stay.”

“And how about you?”

Regretta signed and her eyes grew misty. “I can’t undo the things that put me here.”

“What things? You’ve never told me. And I’ve never heard you talk about them in therapy.”

“Some things are too personal to talk about in therapy.”

“You can talk about them to me,” said Devlin, gently.

“Maybe another time. Will you come and visit me sometimes?”

“Who said I was leaving?”

“You’re cured. We agreed on that.”

“I didn’t agree. I just haven’t gotten around to explaining why not.”

“So why don’t you agree?”

“Because I’m still ignorant. All we’ve done so far is model the iceberg risk that’s above the surface. What about the 90% that lies beneath?

“Above the surface? Below the surface? Devlin, what are you talking about?”

“Sorry, I told you I don’t communicate very well. What I mean is that the models we’ve been looking at are superficial. The returns in the various regimes don’t depend on our beliefs. On reflection that doesn’t make sense. What really happens and what you believe will happen ought to depend on each other.”

“I don’t get it. A bull market has one distribution of returns. A bear market has another distribution of returns, generally with lower mean and higher volatility. What do your beliefs have to do with it?”

“Plenty. Let’s do a thought experiment. Suppose we’re currently in a bull market, which we believe will last forever. Suddenly a mysterious but very authoritative stranger informs all us investors that a bear market is likely to start next week. What happens to asset prices today?”

“They spike down because we have to take into account the likely reduction in future returns. Oh, I see now. The odds of various regimes affect the future expected payoffs, which should be discounted into current prices.

“Good. Now for a second thought experiment. Are we in a bull or bear market right now?”

“That’s a tough one. I mean, we’ve been in a bear market for some time but it looks to me like it’s drawing to a close. Stocks had a runup a few weeks ago but it kind of faltered. Maybe it will resume soon. Yes, maybe that runup was the beginning of a bull market. Well, then again we’ve seen so many false bottoms recently, and recent unemployment figures were very discouraging. So perhaps we’re still in a bear market after all.”

“In other words, you’d feel more comfortable saying that with *X*% probability we’re in a bull market and with 100-*X*% probability we’re in a bear market.”

“Definitely. Actually, if you will allow the possibility of a sideways market I’ll lay odds on that too.”

Devlin smiled and wagged his finger. “That’s cheating. Stick with the thought experiment. Now suppose new economic statistics are released that surprise on the upside. As a result, you and every other investor decide that the chances we’re currently in a bull market are really *X*+1%. What should that do to asset prices?”

“Raise them slightly, of course.”

“In other words, current prices ought to depend on your beliefs about both the nature of the current regime and the likelihood of various future ones. Yet the regime-switching models we built for Conway don’t appear to take that into account.”

“Oh my. How do you propose to fix them?”

“I don’t know,” said Devlin, sadly. “I really don’t know.”

Regretta looked into his eyes. *Poor Devlin.* “It must be lonely to see things you can’t get others to notice.” She reached out and put her hand on his.

Devlin smiled back gratefully. “You noticed. Conway noticed.”

“There will be others, Devlin. Just be patient. It always takes a while for a new approach to win converts. In the meantime, don’t lose sight of what you’ve already accomplished.”

“You mean of finding holes in portfolio theory without filling them?”

“You didn’t find just any hole, Devlin. You found an abyss. And then you figured out how to navigate it.” Regretta saw Devlin start to object. “OK, not perfectly. But what you’ve come up is so much better than the status quo.”

Devlin laughed wryly. “Not if you care less about the real risks than in trussing up their appearances. Which seems to be how most people in finance think.”

“Appearances will always matter in finance. You know that. But real risks matter too. And more people will come around when they appreciate the difference.”

Devlin gazed at Regretta and sat very still. Finally he reached out and touched her arm. Leaning over, he kissed her on the cheek. “Thanks,” he said softly.

“I think we’re attracting attention,” whispered Regretta, looking around. “Perhaps we should continue this conversation elsewhere.”

“I’d like that. I’d like that very much. But I would hate for it to cause regrets later.”

Regretta smiled at Devlin. “There will be no regrets.”

They drifted slowly down the corridor, hand in hand. It was a wonderfully clear evening, and they stopped below a skylight to watch the stars. Then a familiar voice broke the spell. Startled, they turned around.

“Hi, Devlin. Hi, Regretta.” said Conway. “It’s great to see you again. I’m really looking forward to continuing our discussions.”

Devlin instantly let go of Regretta’s hand. “Hi, Conway,” he said, a bit embarrassed. “I didn’t realize they allowed visitors on Friday nights.”

“Who said I was visiting?”

(to be continued).

FURTHER READING

While there's a lot more to Devlin, Conway, and Regretta's story, you won't find it here. Instead I want to tell you about some books and articles that fill in missing pieces or provide a different perspective. I focus on a few I found especially interesting without trying to present a comprehensive list. Still, between the books mentioned and the sources they cite you can find plenty of food for thought. To facilitate your search I have organized the recommendations by topic.

The Magic of Risk

Primitive peoples, much like personal injury lawyers today, tended to reject the notion of blind chance. That's hardly surprising. It's always tempting to divine an underlying intent, and hard to prove the diviner is wrong. But even scientists find it hard to grapple with risk. How do you pin down the concept of something that can't be pinned down?

Eventually people boxed risk into neat theories, which are widely taught today. Unfortunately, seeing risk only inside boxes is like seeing lions only inside cages. To restore your sense of magic and wonder, I heartily recommend Peter Bernstein's bestseller Against the Gods: The Remarkable Story of Risk (New York: John Wiley & Sons, 1998). It conveys excitement about both the centuries of struggles for understanding and the frontiers still left to explore.

One of those frontiers is our own mind. Coping daily with risks for millions of years, our ancestors bequeathed each of us some deep intuition for risks. In the jungles of Africa this intuition helped save our ancestors' skins. Unfortunately in the jungles of Wall Street it often leaves us ripe for skinning. Few people understand those weaknesses better than Nassim Taleb, a successful Wall Street trader with a PhD in statistics. His Fooled by Randomness (New York: TEXERE, 2001) is full of entertaining and insightful tales.

Probability Theory

If your math is so rusty you can't recall what you forgot, I recommend the Schaum's Outline Series (e.g., Probability and Statistics, 2nd Edition (New York: McGraw-Hill, 2000) by Murray Spiegel, John Schiller and Alu Srinivasan) for quick refreshers. I especially like the Schaum approach of presenting each new idea in a separate capsule and have tried to emulate it.

Before I could shave I had the good fortune of studying math at Princeton, home of the eminent probability theorist William Feller. Decades later, having squandered my childhood gifts, I dusted off my old copy of his An Introduction to Probability Theory and Its Applications, volume I, 3rd edition (New York: John Wiley & Sons, 1968) and tried to resuscitate old neurons. A few months later, I added volume II, 2nd edition (1971). If you have the time I highly recommend both. They're not the crispest ticket but they can help teach you how to think…or re-teach you.

Granted, if you stop at Feller you'll miss a lot of nifty tricks. For a concise compendium of these tricks and the underlying derivations see Kenneth Lange's Numerical Analysis for Statisticians (New York: Springer, 1998).

One topic hardly covered in the above is mixed multivariate normality. If that's what you're thirsting for, drink to your heart's content from Finite Mixture Models (New York: John Wiley & Sons, 2000) by Geoffrey McLachlan and David Peel.

Finance Theory

To brush up on basic finance theory I recommend a good business school textbook like Principles of Corporate Finance, 6th edition (New York: McGraw-Hill, 2000) by Richard Brealey and Stewart Myers. It provides perspective that deeper treatments often lack.

Jonathan Ingersoll's Theory of Financial Decision Making (Savage, Md: Rowman & Littlefield, 1987) provides a more rigorous treatment. It's crisp and well-organized, though dry as dust. I use it a lot as a reference.

When you're ready to tackle derivatives pricing, pick up one of Paul Wilmott's books. The latest and greatest is the two-volume Paul Wilmott on Quantitative Finance (Chichester, UK: John Wiley & Sons, 2000). It's hard to imagine a clearer treatment and it's good for a few laughs. Also, check out www.wilmott.com, currently the most bubbling site in quantitative finance.

A lot of cutting-edge finance tends to develop everything through martingales (the concept of a fair game made rigorous). I haven't bothered with them here, for the same reason I don't drag out cannons to swat flies. But if you want to master advanced weaponry, Marek Musiela and Market Rutkowski's Martingale Methods in Financial Modelling (Berlin: Springer-Verlag, 1998) provides good training. For even better training, gear up first with David Williams' Probability with Martingales (Cambridge, UK: Cambridge University Press, 1991).

Risk Management

This book gave short shrift to standard value-at-risk methodology. For a much fuller treatment that is very readable and offers some constructive patches, read Kevin Dowd's Beyond Value at Risk: The New Science of Risk Management. (Chichester, UK: John Wiley & Sons, 1998).

Portfolio managers wanting detailed practical advice should buy Active Portfolio Management (New York: McGraw-Hill, 1999) by Richard Grinold and Ronald Kahn. While it doesn't formally incorporate iceberg risk the way I do, it covers a host of things I don't.

Fischer Black and Robert Litterman wrote up their portfolio model in *Asset Allocation: Combining Investor Views with Market Equilibrium*, a paper distributed in 1990 by Goldman Sachs. See also the 1999 Goldman Sachs paper *The Intuition Behind Black-Litterman Model Portfolios* by Litterman and Guangliang He.

If all this leaves you hungry for more advanced statistics, check out Theory of Financial Risks: From Statistical Physics to Risk Management (Cambridge, UK: Cambridge University Press, 2000) by Jean-Philippe Bouchaud and Marc Potters. By focusing on the characteristic functions of probability distributions it address far more advanced topics than I do, like the risks from hedging options in discrete time. On the negative side, its treatment of regime change is very weak. I hope that future developments will integrate their approach with mine.

Physics

You don't need to know any physics to understand finance theory, but it helps. I especially like books that explain how physics came to embrace uncertainty, because they promote hope that someday finance theory will do the same. The Transnational College of LEX provides a delightful mix of history and math in What is Quantum Mechanics? A Physics Adventure (Boston: Language Research Foundation, 1996), translated from the original Japanese edition of 1991.

Erwin Schrödinger delivered some brilliant lectures in Dublin in 1944 on partition functions and related topics. I found them in a 1989 Dover reprint of Statistical Thermodynamics (Cambridge, UK: Cambridge University Press, 1952). But for maximal inspiration read books by or about Richard Feynman; e.g., his three-volume Lectures on Physics co-authored with Robert Leighton and Matthew Sands (Reading, MA: Addison-Wesley, 1963-1965) or Genius by James Gleick (New York: Random House, 1992). While Feynman couldn't have been Devlin's father, you'll see why I pretended he was.

The Art of Seeing

Sometimes a conformist reality envelops us so tightly that we only see it clearly in our dreams. But to find master dream-weavers we must turn from science to art. Want to see Stalin visit Stalinist Moscow and fit right in? Read The Master and Margerita by Mikhail Bulgakov. Want to see Latin American history endlessly repeat itself in the lives of a single family? Read One Hundred Years of Solitude by Gabriel Garcia Marquez. Want to see the tyranny of well-meaning technocrats? Read Brave New World by Aldous Huxley. I've drawn inspiration from all these sources and smuggled a few allusions into this book.