A guide for the perplexed quant

Quantitative financial modelling seems to employ both the language and techniques of physics, but how similar are the two disciplines? *Emanuel Derman* comments on the practice of financial modelling and the environment in which it is done.

Life for quants has changed since I first came to work on Wall Street in 1985. Then, I quickly noticed the embarrassment involved in being 'quantitative'. Sometimes, when talking in a crowded elevator to another 'quant', you might start to say something about the duration or convexity of a financial instrument. If the person you were talking to had been at the firm a little longer than you, he—it was usually a he—would cringe a little, and rapidly try to change the subject. 'See the Yankees game last night?' he might ask, or 'Futures dropped more than a handle today!' he might exclaim, the sort of things a real bond trader might say. Soon, you began to realize, there was something a little shameful about two consenting adults talking math in a crowded elevator; there was something embarrassing about mentioning UNIX or C in the company of traders, salespeople and bankers. There was something awful about being 'outed' as a quant in public. People in the elevator just looked away.

Even five years ago, quantitative skills were reluctantly accepted. Once, a friend and I were talking on the trading floor when one of the convertible traders walked between us, momentarily. Suddenly he grimaced and winced; he clutched his temples with both hands as though a sharp pain had pierced him and exclaimed, 'Aaarrgghhhh! The force field! It's too intense! Let me get out of the way!'.

In those days, I always used to make a bet with anyone new to our group. Sometime soon, I would wager, some junior trader or salesperson will enter the elevator we're in and, seeing all of us together, will say something like: 'Uh-oh! Isn't there some company rule against all of you being in the elevator at once?'.

Sure enough, I usually won.

The coolest approach to all of this I ever saw was that of Stan Diller, then head of Goldman's Financial Strategies Group in the early 1980s. In an article in Forbes about his group, I recall him being asked what degree he had. 'A PhD', he told the reporter, 'but don't tell the people I work for—they'll knock a half million off my pay!'

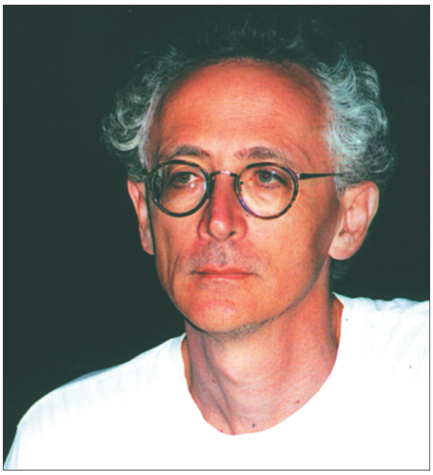

Uh-oh! Isn't there some company rule against all of you being in the elevator at once?

Things have obviously improved in 15 years, partly thanks to the newly found respect for technology as manifested in dot.com stock prices, partly thanks to the mystique surrounding quantitative trading and hedge funds, and partly because of the genuine need for quantitative risk management. The International Association of Financial Engineers itself has played a large role in making quantization a respectable profession rather than just an intersection of various self-taught skills.

Now, I want to get more serious. I want to examine the experience of building models on the Street. I want to reflect on the sociology of quants and traders, on what works in the field of financial modelling, both scientifically and sociologically. I want to present a sort of *guide for the perplexed quant*.

1. Modelling styles

Deep inside, everyone recognizes that the purpose of building models is divination—foretelling the future, and controlling it. Before I tackle the question of what works and what doesn't, I want to analyse the different approaches to modelling. In doing this, I'm going to rely on my upbringing in physics where, over its much longer history, everyone has come to understand that there is more than one way to model the world.

Fundamental models

The first kind of model is what physicists call a fundamental model. A fundamental model attributes events and effects in the world to deep and essential dynamical causes. It consists of a system of postulates and data, used together to draw causal and dynamical inferences about some system. Fundamental models say 'These are the laws of the world'. They try to describe the dynamics of the world in God's terms; they seek eternal truths. As an example, Kepler's laws of planetary motion state that:

- planets move around the sun in elliptical orbits;
- the line from the sun to the planet sweeps out equal areas in equal times;
- the square of the period of a planet is proportional to the cube of its radius.

This is not quite a fundamental model; there's no causality and dynamics. But nevertheless, there is something transcendent in these laws that surpasses mere observation. These are not the observations of a modern data-miner.

There has been a profound distillation of data into structure. Kepler specifies an invariant relation satisfied by the area swept out by a planet's orbital radius. It is, in fact, a restatement of the law of conservation of angular momentum. Also, the laws are stationary, the constants are constant and the exponents are the same for all planets at all times.

Newton added dynamics to this structure. His laws state that:

$$F = ma \text{ and } F = \frac{GMm}{r^2}.$$

These are laws of cause and effect. Force produces acceleration; mass causes gravitational force. This is no longer a model—it's a theory. Newton has isolated the appropriate variables to use and specified a causal relationship between them.

Few financial models have this quality. Black–Scholes–Merton comes by far the closest; it isolates structural variables and describes the dynamics of replication. This is what we should strive for, but it's very difficult and rarely possible.

Phenomenological models

The second type of model is what physicists call a phenomenological model. Phenomenological models, like fundamental models, are built to describe and draw inferences from data, but there is a toy-like quality to their description. They work by a pragmatic analogy that one hopes is descriptive and useful, but one doesn't delude oneself into thinking it's the truth.

Phenomenological models don't say 'This is a law'. Instead, they say 'Approximately, you can think of this part of the world as being a lot like this other kind of thing you already understand'. To be a little pretentious, these are models that describe the world in man's language rather than God's.

An example is the liquid drop model of the nucleus, where you think of an atomic nucleus as rather like an oscillating drop of fluid, even though you know that at a more reductionist level it's composed of other particles. Calibrating the drop's parameters to match some known properties of the nucleus, you can then use the model to compute and predict values of other yet unmeasured properties.

This approach is very similar in spirit to some of the most widely used yield curve models in finance, where simple prescriptions for the future evolution of rates are calibrated by fitting today's bond and cap prices, and then used to value other, more illiquid, instruments.

Statistical models

Third, there are statistical or regression models. Statistical models also consist of data and (statistical) inferences, but dynamics is largely absent. This is the simplest kind of modelling, often useful where practice requires some answer, no matter what.

Physicists and financial modellers both use statistical analysis to find the best estimate for various parameters in a model, or to quantify how well a *dynamical* model's predictions agree with experiment. But only financial modellers use statistics to search for relationships without dynamics. It can be useful, but it's more limited.

2. Some opinions about models

Finance isn't physics

Physics aspires to fundamental models, with predictions of eight decimal places. Most financial models are phenomenological toys based on analogy. Financial modelling is *never* going to provide eight-decimal-place forecasting.

You shouldn't really expect it to. In physics, you're playing against God, and God doesn't change the laws very often. You're trying to describe the created world, with a combination of intuition and experiment and mathematics. You use parameters like mass and charge, which are time-independent and not obviously of human origin.

In finance, you're playing against God's creatures, agents who value assets based on their feelings about the future in general and their future in particular; these feelings are ephemeral, or at best unstable, and fresh news on which they are based keeps streaming in. Finance uses parameters like future risk and future return, which not only refer to the future rather than today, but also are opinions rather than facts. Expected value in finance clearly derives from human beings who are doing the expecting, while mass and charge and electromagnetic force apparently don't—or at least not in an obvious way.

Physics is much more amenable to mathematics and precision. Finance is vaguer and, consequently, not easier but much tougher. Our Newton hasn't appeared yet, and the three wise men haven't even clearly seen the light in the sky. That's why there are many more physicists than quants. It's easier to do useful work within a reliable framework that's already been established.

Similarly, within finance, fixed-income securities, with their multiple, precisely known, cash flows, are much more amenable to mathematics than are equities. In a sense, everything in fixed income is a derivative of the yield curve; in equities, only derivatives are derivatives. That's why there are more fixed-income quants than equities quants.

There's a crisper way of putting this. I once said to one of our equity traders that I thought, on average, fixed-income traders seem to be smarter than equity traders. He agreed, adding 'That's because there's no competitive advantage to being that smart in equity trading'.

Similarly, perhaps there's less of a competitive advantage to being quantitatively smart in finance.

Financial models are interpolation models

According to Steve Ross, '... options pricing is the most successful theory not only in finance, but in all of economics'.

I think that's so because the fundamental problem of options theory is the valuation of hybrid, nonlinear securities and options theory is an ingenious but glorified method of interpolation. That's not an insult. Traders use options theory to convert simple, linear perceptions about thinkable quantities like volatility and probability into complex nonlinear patterns of variation in price. They do this by regarding a hybrid option as a probability-weighted mixture of simpler securities.

Black–Scholes works so well because it's trying to solve a relative-value problem rather than an absolute-value problem. You give it the known prices of a stock and a bond and it tells you the price of the mixture in an option. Physics actually works well on absolute value problems as well.

The real world violates most of the

principles of options theory. Illiquidity and transactions costs mitigate the law of one price. Evolution isn't log-normal. Volatility isn't deterministic. Replication is neither continuous nor costless. Nevertheless, the idea of an option as an interpolating mixture is robust despite all these oversights.

Financial models turn opinions into dollar values

The fundamental problem in financial modelling is how to turn opinions about the future into current dollar values. It's easier to have opinions, even wrong ones, than it is to guess at value. As a trader, all Black–Scholes requires of you is your opinion about future volatility. It then converts your conceptual thoughts about future uncertainty into a fair dollar value. This is no black box or voodoo model; it's reason transmuted to numbers, and that's the right way for a model to work.

Now, people's opinions are just opinions, vague and uncertain. If you are going to build models that get primed by human opinions, many pragmatic consequences follow.

Factors should be things you can opine about

So, better to have market models with variables and factors you can name and whose nature you can grasp and opine about, than to have black-box models that dictate actions without a perceived structure.

Financial models should be simple

If your foundation is an opinion, and therefore necessarily vague, don't build a house of cards on it.

Ravi Dattatreya tried to tell me this in my first week at Goldman. He asked me to enhance a Black–Scholes-style bond options model he had built. I started out slowly and carefully, tackling the problem the way I used to in physics. After about a week, he got impatient, perhaps too soon. 'You know', he said, 'in this job you really need to know only four things: addition, subtraction, multiplication and division—and most of the time you can get by without division!'.

New quants on Wall Street often are amazed at the naivete of Black–Scholes, and immediately try to do better by adding jumps, stochastic volatility, correlations, transactions costs, etc. But traders are limited humans and data are sparse, so this extra complexity doesn't necessarily improve things. A usable model has to provide both input and a way of speaking that comes naturally.

Of course, as time goes by, people can, and do, get more sophisticated, layer by layer, with new theories spawning new strategies which then spawn even newer models.

'You can't give someone a Black–Scholes calculator and turn him into a trader'

Lee Vance, a former Goldman trader, once told me this. Because models turn parameter estimates into prices, the person using the model has to have a visceral feel for the parameters and their dynamics. If this isn't self-evident to you, you haven't been here long enough. Traders have to learn to think in terms of the model's parameters and modellers have to learn to build models that cater to a trader's natural mental framework, or there's no meeting ground.

'What was the world created for?' asked Candide. 'To drive us all mad' said Martin

When I started at Goldman, I was still in the throes of the physicist's seductive dream of a unified theory of everything. I became entranced by the idea that one could build a single model for all interest-rate-sensitive instruments. Fischer Black, who had been around much longer, was sceptical. He always insisted that it was perfectly acceptable to have different models for different domains.

Time and exposure to the real world taught me that he was right. One should be ambitious, but not too ambitious. The world of financial valuation is complex, and what works are simple low-dimensional models with a few essential characteristics. Most real things are too messy for a full theoretical treatment, and that's why implied values, which mask so many unknowns in one effective calibration parameter, play such a large role.

Because of this, financial valuation will always have much in common with art or antiques valuation, where knowledge and experience and street sense are as important as any formula.

There's a simple sort of modelling uncertainty principle that's operative most of the time. The more factors you need to calibrate to, the less useful your model.

Optimization in finance has limited value

In engineering, where the laws are understood, or in travelling-salesman problems, optimization is sensible because each scenario is precisely understood, and you're trying to find the best one. In financial theory, in contrast, each scenario is imprecisely wrong—there's a crude interest rate model, a crude prepayment model, and other misspecifications. While averaging may cancel much of the misspecification, optimization tends to accentuate your lack of knowledge.

3. Quants and traders

Practical problems are your bread and butter

Academics have glorified notions about the use of models because they haven't had to manage the infrastructure that models demand. Using a model for trading needs a portfolio system, a product database, a graphical user interface, live price feeds, calibration. It just needs a few months on the Street to see that the overwhelming practical limitation in making use of models is the state of software systems.

So, don't shy away from practical problems and software. Most of my time over the last ten years was spent earning my bread and butter by developing risk management software. Many of the most interesting problems come from talking to traders about risk and its interface. The nice thing about working on Wall Street as opposed to in academic life is that when you lack inspiration you can do lots of useful things that require only time, not genius.

Quantization and its discontents

Quants are neither traders nor salespeople, programmers or mathematicians. Quants have to be interdisciplinary. One of the theories about what makes an animal non-kosher is that it crosses categories. The Creation in the book of Genesis categorizes animals by both their species and their habitat, referring to 'birds of the sky' and 'fish of the sea'. According to some religious commentators, not belonging exclusively to any particular category is what makes an animal non-kosher. Shrimp live in the sea but aren't fish and don't swim. Ostriches are birds, but don't fly in the sky. Similarly, cloth made from a mixture of linen (a plant) and wool (an animal product) is also proscribed.

It often strikes me that quants are the non-kosher category-violators of investment banks; they need to do a bit of everything to perform well and they provoke discomfort in regular one-category people like traders or salespeople. A quant is an interesting amateur, with no clear professional role model.

In the long run this is bound to improve, but many quants rapidly start aspiring to be something else, something less dimensioned.

Dealing with traders

Quants and traders have fundamentally different temperaments. Quants come from a background where they need to like to do one thing deeply and well, and not stop until they are finished. Work on the Street often needs several quick approximate answers. The hardest adjustment, when I moved to Wall Street, was to learn to do many things in parallel and not too badly, to interrupt one urgent and still incomplete task with another more urgent one, to complete that, and then 'pop the stack'.

Quants are deliberate by nature. Traders have to be opinionated, visceral, fast thinking and decisive, though not necessarily always right. It takes a long time to learn to talk to traders. This isn't helped by the fact that they're always busy and distracted, and it takes an hour of uneasy hovering to have five minutes of punctuated exchanges. If you want to convey information to a trader, you have to learn to start from the conclusion, and you have to learn to be brief. Those are good skills in general.

These temperamental differences are most apparent when our firm does its annual 360-degree performance reviews, in which everyone's strong and weak points are noted by all their associates. Generally, a quant reviewing another quant lists all of his good qualities in great detail—the technological knowledge, the mathematical skill, the ability to be inventive and so on. Since many quants share these qualities, the reviews of quants by quants tend to lack differentiation.

Traders' reviews are more likely to get to the point. One trader I know wrote 'this guy is smart but he talks too much, and we don't have time to listen. Tell him it puts people off'. Traders and quants tend to think differently about financial value. Traders need to communicate so incisively because the essence of good trading is responding intelligently to variation (and the threat of variation). Derivatives' prices, being nonlinear, magnify variation. As a consequence, derivatives traders think naturally about change. When they consider an option, they think about what may happen to its price tomorrow as markets move. Traders think about scenario hedges. Quants think more about current value and how to compute it as an expectation over all scenarios.

> Shrimp live in the sea but aren't fish and don't swim. Ostriches are birds, but don't fly in the sky

Feynman-Kaç tells you that the two approaches are complementary. Nevertheless, this difference in viewpoints—the quant's static view of all future scenarios versus the trader's dynamic look at tomorrow's changes—tends to make communication uneasy. The more you can learn to think both ways, the better you can communicate and understand.

Options traders can get by with less math than you think. Tour-de-France cyclists don't need to know how to solve Newton's laws in order to bank around a curve. Indeed, thinking too much about physics while riding or playing tennis may prove a hindrance. But good traders do have to have the patience to understand the essential mechanism of replicating the factors they're trading.

Traders like to be anecdotal. They see movements as responses to supply and demand. Quants like to be structural; they describe the same movements in terms of models. These don't have to be competing pictures; they can be complementary descriptions of the same thing.

I like to think that what the trading world needs most is people who understand both models and markets. But even that is not a panacea; the people at LTCM were about as well rounded in this respect as anyone.

4. Generalities

Many quant resumés from headhunters recite the litany 'Knows HJM, knows BGM, knows GARCH ...'. But, in research, you don't just want someone who knows a lot. You also want someone who can persevere, who can spend months or even years building a model, someone who can dig himself out of a hole he fell into along the way. For this, a PhD is a good, if painful, training.

You can develop the world's cleverest model, but intelligent traders won't embrace it unless they understand it. Nor should they. Therefore, you cannot spend enough time trying to explain your results in qualitative terms. Doing this helps your understanding too.

Traders and salespeople don't always know what they need; they're too involved in the hurly-burly. Don't always wait for someone to commission something from you. Value often arises out of one person's inspiration. That's certainly what happened with Black–Scholes–Merton. No trader in the world, at that time, could have had the imagination to commission anything like that. There's a line from Blake that expresses this well,

> *Improvement makes straight roads; but the crooked roads without improvement are the roads of genius.*

We can't all be Blake or Black–Scholes–Merton, but, in your own way, don't be scared to strike out sometimes. Other people don't always know what's best for them.

Research

I think we're going through a relatively fallow support period for quantitative research on the Street right now. The current demand is for electronic distribution and its infrastructure. These things are cyclical.

But there's lots of value to research. A vocal research organization can help spread the culture of wise model use. It can attract business by helping clients solve their investment problems. It can bring more diversely talented people to the firm. All of this is going to become more important as investment banks broaden and their capabilities enlarge.

In the long run, the argument for research and education is a lot like the argument for the North American Free Trade Association. Research it isn't a zero-sum game; it creates value. As more people get comfortable with new concepts, the total pie grows.

5. *Gedanken* experiments: using models appropriately

How do you know when a model is right? You cannot easily prove models 'right' by observation; data (especially accurate data) are scarce, and the financial world of human action and reaction is non-stationary. What's right in one regime is wrong in the next. And worse, financial models are calibrated to our implied view of the future, so you're always testing not only the naked model's dynamics but our views as well. This is much harder than testing whether a physics model is right; physics has no implied human view, only observed quantities—though of course, if you think about it too hard, even the mass and charge of objects are human inventions whose values are extracted from comparing data with theory.

As a physicist, when you propose a new model of the physical world, you're pretending you can guess the structure God created. It sounds eminently plausible. Every physicist believes he has a small chance of guessing right, or else he would not be in the field. But as a financial theorist, when you propose a new financial model, you're pretending you can guess the structure of another person's mind. When you try out a simple yield-curve model, you're implicitly saying something like 'Let's pretend that people care only about future short rates, and that people expect them to be distributed log-normally'. As you say that to yourself, if you're honest, your heart sinks. You know immediately that there is no chance you are truly right.

Fischer Black said that:

> *In the end, a theory is accepted not because it is confirmed by conventional empirical tests, but because researchers persuade one another that the theory is correct and relevant.*

This is painful for some people to accept, especially people who don't spend their time depending on models, but it's largely true. I'd go even farther. From the viewpoint of someone who works with traders, I like to think of financial models as analogous to the way quantum and relativity physicists in the early part of the last century used *gedanken* experiments. *Gedanken* experiments—German for thought experiments—are imaginary experiments, a sort of mental stress-testing of the physical world, done in your head because it was too hard to do in practice, in order to force your theoretical picture of the world into a contradiction. Einstein imagined what he would see sitting on the edge of a moving light beam in order to get insight into the contradiction between electromagnetic theory and Newtonian observers. Schrödinger imagined an unobserved cat sealed in a box with a radioactive atom that, on decaying, would trigger a Geiger counter to release cat poison.

I think that's the right way to use mathematical models in finance. They are only models, not the thing in itself. Don't expect them to be truly right. Regard models, instead, as a collection of parallel thought universes you can explore. Each universe should be consistent, but the real financial and human world is going to be much more complex than any of them. You're always trying to shoehorn the real world into one of them to see how useful that approximation is.

You must always ask: does the model give us a set of plausible variables to use in describing the world, and a set of relationships between them that we can use to analyse and perturb the world? You're always trying to make a limited approximation of reality, with perceptual variables you can think about, so that you can say to yourself, or your boss, 'I was short emerging-market volatility, so we lost money'.

Good theories, like Black–Scholes, provide a theoretical laboratory in which you can explore the likely effect of possible causes. They give you a common language with which to quantify and communicate your feelings about value. Along the lines of Andy Warhol, I'm tempted to summarize this by saying that: 'In the future all models will be right, but each one only for 15 minutes'.

The right way to proceed, then, is to push a model as far as you can, but always be aware of its limitations, of the fundamental assumption, underlying everything you do, that the human world is indeed modellable.

I once read a biography of Goethe, one of the last people to make contributions to both art and science. Scientists regard Goethe as a poet who strayed beyond his proper place. His critics said he mistakenly thought of nature as a work of art, and that he was trying to be qualitative where he should be quantitative. But, according to the book I read, Goethe was not so naive as to think that nature is a work of art. Rather, he believed that our knowledge and description of nature is a work of art.

That's how I like to picture what we do in financial modelling—making a beautiful and truthful description of what we can see. We're involved in intuiting, inventing or concocting approximate laws and patterns. We can synthesize both art and science in creating understanding. We can use our intuition, our scientific knowledge and our pedagogical skills to paint a picture of how to think qualitatively, and then, within limits, quantitatively, about the world of human affairs, and in so doing, have an impact on how other people think.

The success of options valuation is the story of a simple, asymptotically correct idea, taken more seriously than it deserved and then used extravagantly, with hubris, as a crutch to human thinking. This always reminds me of an aphorism from Blake's 'Proverbs of Hell' in *The Marriage of Heaven and Hell*,

> *If a fool would persist in his folly, he would become wise.*

I think that's what we've done with options theory. But the catastrophes of options valuation are the obverse side of the same coin, when people pay more attention to formulae than ideas, so that extravagance evolves into idolatry. Blake has an aphorism for this too, in his 'Proverbs Of Hell',

> *Bring out number, weight & measure in a year of dearth.*

Somewhere between these two extremes, north of hubris but still south of idolatry, lies the wise use of models. It always takes human judgement to draw the line.

Emanuel Derman is a Managing Director in Firmwide Risk at Goldman, Sachs & Co. in New York.

This article is adapted from a speech Emanuel Derman delivered on 12 October 2000, on receiving the IAFE/SunGard 2000 Financial Engineer of the Year Award. A version of this speech was previously printed in 2000 *J. Derivatives* **8** 58–64.