

Quantitative Strategies Research Notes

April 1996

Model Risk

Emanuel Derman

Copyright 1996 Goldman, Sachs & Co. All rights reserved.

This material is for your private information, and we are not soliciting any action based upon it. This report is not to be construed as an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. Certain transactions, including those involving futures, options and high yield securities, give rise to substantial risk and are not suitable for all investors. Opinions expressed are our present opinions only. The material is based upon information that we consider reliable, but we do not represent that it is accurate or complete, and it should not be relied upon as such. We, our affiliates, or persons involved in the preparation or issuance of this material, may from time to time have long or short positions and buy or sell securities, futures or options identical with or related to those mentioned herein.

This material has been issued by Goldman, Sachs & Co. and/or one of its affiliates and has been approved by Goldman Sachs International, regulated by The Securities and Futures Authority, in connection with its distribution in the United Kingdom and by Goldman Sachs Canada in connection with its distribution in Canada. This material is distributed in Hong Kong by Goldman Sachs (Asia) L.L.C., and in Japan by Goldman Sachs (Japan) Ltd. This material is not for distribution to private customers, as defined by the rules of The Securities and Futures Authority in the United Kingdom, and any investments including any convertible bonds or derivatives mentioned in this material will not be made available by us to any such private customer. Neither Goldman, Sachs & Co. nor its representative in Seoul, Korea is licensed to engage in securities business in the Republic of Korea. Goldman Sachs International or its affiliates may have acted upon or used this research prior to or immediately following its publication. Foreign currency denominated securities are subject to fluctuations in exchange rates that could have an adverse effect on the value or price of or income derived from the investment. Further information on any of the securities mentioned in this material may be obtained upon request and for this purpose persons in Italy should contact Goldman Sachs S.I.M. S.p.A. in Milan, or at its London branch office at 133 Fleet Street, and persons in Hong Kong should contact Goldman Sachs Asia L.L.C. at 3 Garden Road. Unless governing law permits otherwise, you must contact a Goldman Sachs entity in your home jurisdiction if you want to use our services in effecting a transaction in the securities mentioned in this material.

Note: Options are not suitable for all investors. Please ensure that you have read and understood the current options disclosure document before entering into any options transactions.

SUMMARY

This report analyzes the assumptions and risks involved in using models to value financial securities.

Emanuel Derman (212) 902-0129

I am grateful to Alex Bergier, Barbara Dunn,
Didi Hu and Iraj Kani for comments.

INTRODUCTION

Securities markets in the past twenty years have seen the emergence of an astonishingly theoretical approach to valuation, market-making and arbitrage in complex market sectors. Many securities firms now base their bid and offer prices for complex securities on detailed analytic or computer models built by scientists¹. Most of this theory centers around derivatives, instruments whose value stems from their contractually defined relation to more elementary securities or market parameters. In this generalized sense derivatives encompass many products: index futures and options are derivatives on the underlying index, CMOs are derivatives on interest and prepayment rates, and we can even regard bonds as derivatives on interest rates. There are many more examples, from convertible bonds to credit derivatives.

Theoretical models abound. In the fixed-income world, the theoretical approach was probably sparked by the shock to bond portfolio values as interest rates jumped in the late 70's. Duration, convexity, and other theoretical risk and sensitivity measures grew in both sophistication and popularity. You can now attend two-day courses in fitting yield curves and extracting zero-coupon rates. The increased interest-rate volatility also triggered the development of caps, floors, swaps and swaptions, whose valuation and trading were all heavily model-driven. In the equity world, program trading off the mismatch between actual futures prices and their theoretical fair value was made possible by rapid electronic computation and trading.

Equity and fixed-income option trading and structuring grew in part because of the confidence that developed in using the Black-Scholes model and its extensions. The growth in model building and model adoption has also depended on the rapid acceleration in computing power. Computing and modeling have played a sort of leapfrog: more power allowed for fancier models which then ran too slowly, and so in turn required even more power. Advanced users now think of hedging exotic equity index options with standard options, so that one man's derivative has become another man's underlying².

This reliance on models to handle risk carries its own risks. In this report we analyze the assumptions made in using models to value securities, and list the consequent risks.

-
1. Or ex-scientists, depending on your opinion about what they do.
 2. This is not as unprecedented as it sounds. Money itself is a derivative that gets its value from its convertibility into more consumable assets.

WHAT ARE MODELS?

There are at least three different meanings implied by the word *model* in finance, namely:

1. A fundamental model: a system of postulates and data, together with a means of drawing dynamical inferences from them;
2. A phenomenological model: a description or analogy to help visualize something that cannot be directly observed; and
3. A statistical model: a regression or best-fit between different data sets.

Most common financial models fall predominantly into one of these categories.

Fundamental models cover models like the Black-Scholes theory, in which a set of postulates about the evolution of stock prices, data about dividend yield and volatility, and a theory of dynamical hedging together allow the derivation of a differential equation for calculating options values. These are models that attempt to build a fundamental description of some instrument or phenomenon.

Phenomenological models are less fundamental and more expedient, but may be equally useful. For example, some simple bond option models treat the yield of the underlying bond as being normally distributed. This is a useful picture with a plausible feel to it. But it's only a toy, good in a limited range, and not as deep or insightful a description as the Black-Scholes model.

The first two classes of models embody some sort of cause and effect. The last class, statistical models, rely on correlation rather than causation. Users of these models probably hope that the correlation is a consequence of some dynamics whose detailed modeling they are avoiding or postponing. An example is a mortgage prepayment model that regresses prepayment rates against various long- and short-term interest rates and mortgage lifetimes. Modelers imagine homeowners performing certain cost-benefit analyses in deciding whether and when to prepay. Strictly, statistical models describe tendencies rather than dynamics. But knowing tendencies, if they really exist and persist, can be valuable.

SOME FACTS ABOUT MODELS

Models assume cause and effect

Financial models' variables may be people's opinions

Models translate opinions into values

Uncertainty is fundamental

When you build a valuation model of any type, you are implicitly assuming that the objects of your concern are *causally related* to each other, and that the relationship is *stable*, at least for the time that you intend to apply the model.

In the physical sciences where quantitative modeling originated, the variables in models are universal quantities like time, position and mass, that (presumptively) have an existence even when human beings are absent. In contrast, in the financial world, you are dealing with variables that clearly represent human expectations. Even the simplest statement "More risk, more return" refers to *expected* risk and *expected* return, not realized quantities. These are hidden variables: they cannot be directly observed except perhaps by surveying market participants, or by implying their values insofar as they impact other measurable quantities *by way of* a theory or model. Thus, models that use concepts like return or volatility are in most cases assuming a causal and stable connection between the values of these hidden (often unarticulated) variables and security values.

You can start to see how many links there are in the chain from model to usage.

Model users don't just switch on a model and trade according to its results. Having a valuation model doesn't absolve the model user from thinking about the value of a security. Instead, it makes the security value a dependent variable, and requires the user to think about and estimate the values of other independent variables that are easier to grasp and quantify. Mostly, a security valuation model is a way of translating one's thoughts and intuitions about these other variables into a dollar value for the security. For example, the Black-Scholes options valuation model asks a user for an estimate of future volatility, and then translates that estimate into a fair option value. Variations in volatility are much smoother and less dramatic than variations in option value. In this way, good models make it easier to extrapolate security values known under a limited range of market conditions to more distant regimes.

The overwhelming unknown in financial models is certainty. In the physical sciences, the mathematics of statistics and distributions and finally, the calculus of stochastic processes, made their appearance late in the drama. In the financial world, they are the first actors on stage. Everyone expects to predict the position of a man-made satel-

lite, let alone Newton's falling apple, with high precision. No one expects to predict the value of a stock in the future with much precision at all.

Models need domain knowledge

The financial domain is a nitty-gritty world filled with stocks that trade only at certain times with discrete ticks. Usable models exist for some particular sector with particular trading rules, settlement conventions, and other practicalities. Models and modelers need intimate knowledge of the domain they are working in. Financial modeling is as much about content as it is about technical skills.

Financial models are software

Financial models most often end up being implemented as computer programs, either because they need to do many simple things rapidly and repeatedly, or because they need to draw on large amounts of stored information, or because no simple analytic solution to the mathematics is available, and so numerical techniques are required. In addition, much of the gain from using models comes from applying them to portfolios of securities. The handling of portfolios on a computer requires the construction of databases, user interfaces and price feeds. So, both the model itself and the mechanism for employing it involve building software.

A model is only a model...

The real world is often an inchoate swirl of actions, occurrences, facts and figures. There are more things than we've even thought of naming or categorizing. So, even the finest model *is* only a model of the phenomena, and not the real thing. A model is just a toy, though occasionally a very good one, in which case people call it a theory. A good scientific toy can't do everything, and shouldn't even try to be totally realistic. It should represent as naturally as possible the most essential variables of the system, and the relationships between them, and allow the investigation of cause and effect. A good toy doesn't reproduce every feature of the real object; instead, it illustrates for its intended audience the qualities of the original object most important to them. A child's toy train makes noises and flashes lights; an adult's might contains a working miniature steam engine. Similarly, good models should aim to do only a few important things well.

CONSTRUCTING
MODELS

You can understand the things that go wrong with models if you understand how they are developed. Model building is as much art and apprenticeship as engineering and science. Nevertheless, it's possible to delineate some of the procedures involved in constructing a financial valuation model:

- Understand the securities, the markets and the way market participants think about valuation and risk factors.
- Isolate the most important variables that participants use to analyze value and risk.
- Decide which of these variables are susceptible to mathematical modeling.
- Separate the dependent variables from the independent variables. Also decide which are directly measurable and which are more in the nature of human expectations, and therefore only indirectly measurable.
- For some variables, the uncertainty in their future value has little effect on security values³, and they can be treated as known to a good approximation. For other variables, uncertainty is critical. Specify the variables that can be treated as deterministic and those that must be regarded as stochastic.
- Develop a qualitative picture that represents how the independent variables affect the dependent ones.
- Think about how to get the market values of independent observable variables, and how to deduce the implied values of indirectly measurable ones.
- Formulate the picture mathematically. Decide what stochastic process best describes the evolution of the independent stochastic variables.
- Consider the difficulties of solving the model, and then perhaps simplify it to make the solution as easy as possible. But only reluctantly give up content for the sake of an easy or elegant analytical solution.
- Develop a scheme for analytic or numerical solution.
- Program the model.
- Test it.
- Embed it in the software and human environment.

3. For example, in valuing stock options, the future uncertainty in interest rates is largely unimportant, because option value varies smoothly with rate and so the uncertainty averages out. There is consequently no need to know the volatility of interest rates. This is not the case with bond options, whose payoffs vary sharply and non-linearly with interest rates.

THE TYPES OF MODEL
RISKInapplicability of
modeling

The most fundamental of risks is that modeling is just not applicable. For example, it's possible that forecasting stock price movements is more like forecasting political occurrences than like projecting spacecraft trajectories, with psychology and gamesmanship more relevant than mathematics. There's always a temptation to think that complex mathematics has an applicability of its own, but you need a vision of how things work and interconnect before you use mathematics to represent it. You need the analogy or picture first; mathematics is largely the language you represent it in.

In terms of risk control, you're worse off thinking you have a model and relying on it than in simply realizing there isn't one.

Incorrect model

At some level, all models are ultimately incorrect. But even without being perfectionist, here are some of the ways in which model development can go wrong:

- You may not have taken into account all the factors that affect valuation. For example, you may have assumed a one-factor model of interest rates. This is probably a reasonable approximation for valuing Treasury bonds, but much less reasonable for valuing options on the slope of the yield curve.
- You may have incorrectly assumed certain stochastic variables can be approximated as deterministic (see footnote 3).
- You may have assumed incorrect dynamics for a factor. For example, you might have modeled bond *prices* as normally distributed for the sake of analytic simplicity. In practice, bond *yields* are more likely to be lognormal. This discrepancy is worse for short-maturity bonds, but may be forgivable for long maturities.
- You may have made incorrect assumptions about relationships. For example, you may have ignored the correlation between corporate credit spreads and corporate stock prices in valuing convertible bonds. Is this correlation important for the particular property of convertible bonds you are interested in extracting from your model?
- The model you developed may be inappropriate under current market conditions, or some of its assumptions may have become invalid. For example, interest rate volatility is relatively unimportant in currency option pricing at low interest rate volatilities, but may become critical during exchange rate crises.
- A model may be correct in an idealized world (with no trading costs, say), but incorrect or approximate when realities (like market frictions) are taken into account.

- A model may be “correct in principle” but the market may disagree in the short run. This is really another way of saying the model is limited, in the sense that it didn’t take other short-term factors into account (including market sentiment) which can influence price.
- A model may be correct, but the data driving it (rates, volatilities, correlations, spreads, and so on) may be badly estimated.
- A model may be reasonable, but the world itself may be unstable. What’s a good model today may be inappropriate tomorrow. For example, the sentiment about interest rates may be linked to gold prices one year, and to oil prices the next.

Correct model,
incorrect solution

You can make a technical mistake in finding the analytic solution to a model. This can happen through subtlety or carelessness. There are some well known published errors or misunderstandings in the case of some complex derivatives, leading to so-called model arbitrage. It takes careful testing to ensure that an analytic solution behaves consistently for all reasonable market parameters.

Correct model,
inappropriate use

There are always implicit assumptions behind a model and its solution method. But human beings have limited foresight and great imagination, so that, inevitably, a model will be used in ways its creator never intended. This is especially true in trading environments, where not enough time can be spent on making interfaces fail-safe, but it’s also a matter of principle: you just cannot foresee everything. So, even a “correct” model, “correctly” solved, can lead to problems. The more complex the model, the greater this possibility.

As an example, most Monte Carlo valuation models require the choice of a number of simulation paths and steps. Speed requires few simulations, while accuracy demands many. Different securities require different simulation parameters to get a reasonable answer. A user who values a high-variance security with the same parameters as a low-variance security can get inaccurate and even biased results.

The only practical defense is to have informed and patient users who clearly comprehend both the model and the method of solution and, even more important, understand what can go wrong. In the above example, one should start by valuing the security with a variety of

simulation parameters, and perhaps more than one solution method, to examine the accuracy and convergence of the results⁴.

Badly approximated solution

You may have errors in the numerical solution to a correctly formulated problem, or there may simply be natural limits to the accuracy of some approximation scheme. Finite difference solution methods can be unstable, inaccurate or converge slowly. Only careful and knowledgeable testing can help here.

Software and hardware bugs

Many of the worst risks center around implementation. These days, models are sophisticated programs, thousands of lines long, with rich data structures that are used to perform detailed computation. Models undergo revisions by people who were not the original authors. Equally important in making them useful, models need user interfaces, position databases, trade entry screens and electronic price feeds. Programming mistakes in any of these areas can lead to widespread and hard-to-detect errors. You can make errors in logic, rounding, counting the days between dates or the coupons to maturity, to name only a few possibilities. In addition there are occasional hardware flaws, like the widely publicized Pentium floating point error.

Similarly, as programmers strive for greater execution speed, the model is at risk from the natural tension between clarity of style and code optimization.

Unstable data

Many models need the future value of some volatility or correlation. This value is often based on historical data. But history may not provide a good estimate of future value, and historical values may themselves be unstable and vary strongly with the sampling period.

4. The following verbatim quote from someone building a model conveys a sense of the conflict involved in releasing it to users:

“It’s always a dilemma to release a (model)... If I do not release it, and tell people to contact me to price...options... people think I am holding back. When I tell people they should be very careful in choosing methods and parameters, they always say ‘I know, I know’ and get a little impatient. I guess one just has to put some trust in those people who use them.”

AVOIDING MODEL RISK

There is no magic strategy for avoiding risk, but the following general guidelines based on experience in our group at Goldman Sachs may be helpful.

Regard models as interdisciplinary endeavors

Models are generally not back-of-the-envelope formulas handed over to “coders” to turn into executable instructions. Modeling is multidisciplinary: it touches on the practicality of doing business, on financial theory, on mathematical modeling and computer science, on computer implementation and on the construction of user interfaces. Models end up as computational computer programs embedded in human and machine interfaces that are themselves computer programs. The risks lie in the knowledge of the business, the applicability of the financial model, the mathematics and numerical analysis used to solve it, the computer science used to implement and present it, and in the transmission of information and knowledge accurately from one part of the model, in the larger sense of the word, to the next. It helps to be knowledgeable in all of these areas in order to notice an error and then diagnose it.

But, in many firms, model users are traders, salespeople or capital markets personnel who may be physically and organizationally removed from the model creators. Furthermore, the model implementors are programmers who are often similarly separated from the model theorists. To avoid risk, it’s important to have modelers, programmers and users who all work closely together, understand each other’s domains well enough to know what constitutes a warning symptom, and have a good strategy for testing a model and its limits. Too much specialization is harmful. In our group, the modelers themselves write production code for insertion into risk-management systems. Programmers and modelers work in closely-knit teams around a particular product or business area. Informed model users are particularly invaluable.

Because of the large role of computing, we also try to accentuate the importance of software engineering as a discipline

Test complex models in simple cases first

Test models against simple known solutions. If you can solve a model in some simple case, by constructing a tree diagram or solving some equation, compare your computer solution to the simple solution and make sure they’re exactly identical.

Test the model's boundaries

Often, a new model overlaps on older and simpler models. In that case, test the boundaries. If it's an option model, make sure that when the option is deep in the money it behaves like a forward. For a convertible bond model, guarantee that it behaves like a straight bond when it's deep out of the money. Too many complex models go wrong because complexity obscured the error in the simple part of the model. One of the most avoidable mistakes I saw was a convertible bond model that innovatively priced many of the options features embedded in convertible bonds, but sometimes counted the number of coupons to expiration incorrectly.

Don't ignore small discrepancies

If there are any small discrepancies noticed by users or programmers, don't ignore them. Track down their origin. Small disagreements often serve as warnings of potentially large disagreements and errors under other scenarios.

Provide a good user interface

Thorough testing is easier with a flexible and friendly interface. We spend much time building interfaces that allow what-if analysis and graphical display of the results of a model under many different scenarios. Even after many years of use, some errors only become apparent when you notice kinks in a graphical display of the model's results.

Diffuse the model slowly outwards

It is impossible to avoid errors during model development, especially when they are created under trading floor duress. Therefore, in addition to being careful, it's important to have an orderly procedure for disseminating the use of a model. So, after the model is built, the developer tests it extensively. Thereafter, other developers "play" with it too. Next, traders who depend on the model for pricing and hedging use it. Finally, it's released to salespeople. After a suitably long period during which most wrinkles are ironed out, it's given to appropriate clients. This slow diffusion helps eliminate many risks, slowly but steadily.

Pride of ownership

One of the best defenses against modeling error is to ensure that both models and systems are built by people who like doing it, and who take pride in their work.

SELECTED QUANTITATIVE STRATEGIES PUBLICATIONS

June 1990	<i>Understanding Guaranteed Exchange-Rate Contracts in Foreign Stock Investments</i> Emanuel Derman, Piotr Karasinski and Jeffrey S. Wecker
July 1991	<i>Valuing Index Options When Markets Can Jump</i> Emanuel Derman, Alex Bergier and Iraj Kani
January 1992	<i>Valuing and Hedging Outperformance Options</i> Emanuel Derman
June 1993	<i>The Ins and Outs of Barrier Options</i> Emanuel Derman and Iraj Kani
January 1994	<i>The Volatility Smile and Its Implied Tree</i> Emanuel Derman and Iraj Kani
May 1994	<i>Static Options Replication</i> Emanuel Derman, Deniz Ergener and Iraj Kani
May 1995	<i>Enhanced Numerical Methods for Options with Barriers</i> Emanuel Derman, Iraj Kani, Deniz Ergerer and Indrajit Bardhan
December 1995	<i>The Local Volatility Surface: Unlocking the Information in Index Option Prices</i> Emanuel Derman, Iraj Kani and Joseph Z. Zou
February 1996	<i>Implied Trinomial Trees of the Volatility Smile</i> Emanuel Derman, Iraj Kani and Neil Chriss