ROLF W. BANZ

Editorial: Zero Correlation between Theory and Practice in Financial Economics?

Introduction

It is one of the primary objectives of the Society that publishes this journal to promote the dialogue between researchers and practitioners. This article deals with some of the less successful aspects of that dialogue. As a practitioner who remembers some of the basics of the classic models, I am sometimes amazed to see how some of the concepts based on those models are being taught or described in textbooks. I am also amazed to see how casually the concepts are being applied. Of course, I realise that financial economics is not the only field in which theory and practice differ a great deal. But nevertheless I feel that some of those issues merit revisiting. The following discussion uses a number of examples. It is not a very scientific piece (no formulae, just a few tables and graphs). It is an attempt to provide some food for thought about the various ways in which we use and abuse the results of scientific research in our field.

*Adapted from a speech given at the St. Gallen Conference of the Swiss Society for Financial Market Research, 10 October 1997. I am grateful to Daniel Wydler for helpful comments.

Problems with "Modern Finance"

What has become known as "modern finance" has in some respects achieved a level of acceptance that is truly amazing. This is particularly true when we look at activities like trading and financial engineering where arbitrage ideas that were developed in the derivatives literature are now universally applied. At first sight, the same appears to be true in areas such as asset pricing and market equilibrium. But unfortunately a closer look reveals that many of the classic results in those areas are misunderstood and/or misapplied by practitioners.

Example: Sharpe Ratio

Before turning to the question of why this should be the case, let me start with an example of what I mean. It is a simple, familiar and apparently fairly harmless example: the "Sharpe ratio". It is, as we all know, defined as the ratio of excess return and standard deviation. It was first introduced in the sixties. Now, it is widely used as a summary measure of the risk-return trade-off achieved by a managed portfolio and appears increasingly in consultant questionnaires, mutual fund ratings and similar contexts. Its purpose is to rank portfolios.

It is, of course, very easy to calculate. But, as has been pointed out in the literature, it does not really allow an unambiguous ranking of portfolios. There are three kinds of problems with the application of the Sharpe ratio:

- It should be applied exclusively to an investor's entire portfolio, not to portions of the portfolio. Therefore, for manager selection purposes, it is utterly irrelevant whether one Pacific Basin equity manager has a greater Sharpe ratio than another. Its use can actually lead to the wrong decision. What matters in the context of picking a specialist manager is the impact of the manager's portfolio on the risk of the overall portfolio. This is shown in the example in Figure 1. Portfolio A has the bigger Sharpe ratio, but the Sharpe ratio of the combination of B and the existing portfolio is bigger. Thus, portfolio B is a better choice in the context of this investor's overall portfolio.
- But even applied to the overall portfolio, the Sharpe ratio does not give unambiguous rankings as shown in Figure 2. Portfolio A has the higher Sharpe ratio, yet some (perfectly rational) investors might prefer portfolio B. Everybody would of course reject portfolio C.
- A third, perhaps more subtle, problem concerns the application of statistical tools to managed portfolios. The Sharpe ratio is typically estimated on the basis of historical returns. Returns of managed portfolios reflect the decisions of the manager. If a manager engages in market timing or even just occasional portfolio revision or uses derivatives or currency hedges, his performance series is unlikely to be stationary or meet the other requirements for the application of the standard statistical toolbox. In addition, we have to convince ourselves that the past performance is for once a guide to future performance.

In my experience, the Sharpe ratio is almost always misapplied or mindlessly applied. It is not really such a harmless exercise. Its use can and does lead to wrong strategic decisions about investment policy and manager selection. My com-

Figure 1: Sharpe Ratio I

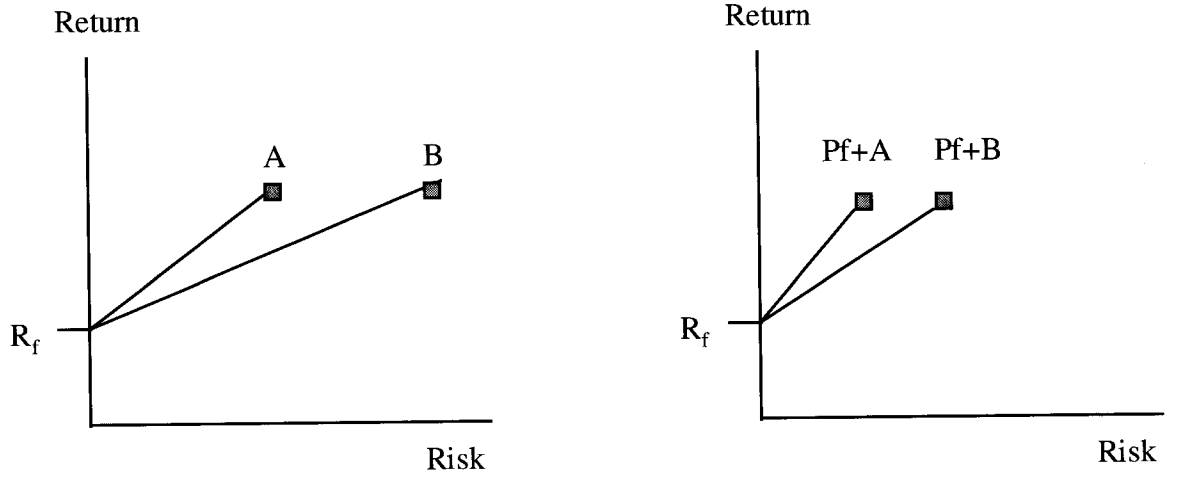

Figure 2: Sharpe Ratio II

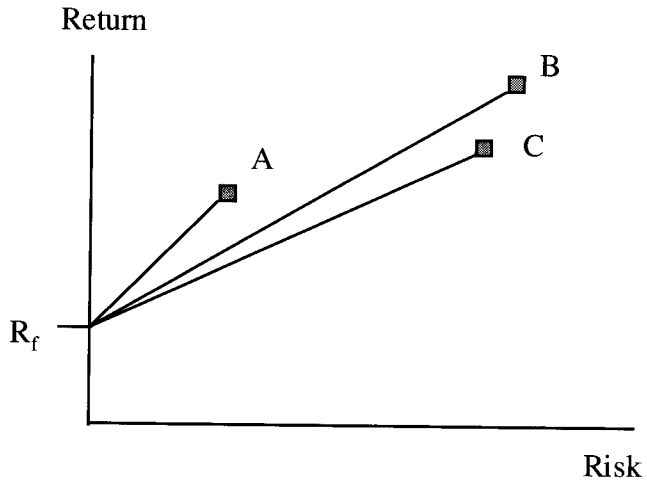

ments on the Sharpe ratio are neither new nor very original. The problems that arise from its misapplication are well documented in the literature. In addition, alternatives are available. Nevertheless, it continues to be taught (after all, it is named after a Nobel laureate) but it is clearly not terribly well explained.

Reasons for Problem

Next, I would like to speculate on the possible reasons why we allow this problem to develop. It appears to me that there are three main contributory factors.

First, there is the process of "simplification terrible" that takes place when a major piece of original research goes on its journey. From original article, it passes through doctoral seminar notes, advanced textbooks, class notes and a few more steps until it ends in a practitioner-oriented "how-to" book. In this process, much of the substance is stripped away. Footnotes tend to go first, then subtle assumptions, caveats and limitations follow until very little other than the basic result is left. What is left is often just the "emperor without clothes".

This is quite understandable: human nature being what it is, we sometimes prefer to learn from a simple source and do not bother to check the original, even though we may feel that the story presented in the textbook sounds a bit too simple.

The consequences of ignoring the fine print can be quite dramatic: Portfolio Insurance (dynamic option replication) was sold 10 years ago as a direct application of BLACK-SCHOLES style option pricing to provide protection from downside risk. Unfortunately, nobody mentioned to the clients that for dynamic replication strategies to work, the assumption of continuous markets is crucial. When the crash came and the New York market opened several per cent below its previous close, that drop was, to the clients' horror, not insured. Yet the result is quite obvious once you think about the basic mechanism of dynamic option replication. As an aside, I wonder whether the typical client of today's guaranteed products whose floor is only an "indicated floor" is aware of what that really means. Of course, it might not help sales if the client was told that what it really means is that

when he actually needs the guarantee badly (after a sudden, big drop), it will be unavailable.
The second contributory factor was illustrated to me earlier this year. A young colleague who is studying for a professional exam (which shall remain anonymous) asked me a question about a detail in a problem in a previous exam. As we discussed the question, I became very impressed by his level of technical expertise. The problem dealt with a swap and he had been taught to do it in continuous time, with semi-annual and annual compounding, with any of dozens of bond pricing conventions and to six decimal places. His question and the discussion showed, however, that in the preparatory courses for the exam, little attention appeared to have been paid to understanding the basic economic principles underlying these techniques. We should not be surprised. It is, of course, much easier to teach and test technical skills than understanding. However, I would be very sorry if we were to breed a new generation of practitioners with impressive titles and with an equally impressive toolbox of techniques at their disposal but little intuitive understanding of the "big picture" and of the underlying assumptions and their consequences.
The third contributory factor is the fact that academics for a long time felt that they had to act as missionaries to turn the heathen practitioners into believers. For that missionary work, the models had to be simple and intuitive and, if possible, supported by some fairly obvious evidence. This led to a serious overselling of the results. This is why some early "modern" finance textbooks sound like excerpts from the bible. But we have come a long way from the days of antagonism between academics and industry that existed in the seventies and eighties. It is time for a more balanced discussion of the issues.

Examples

You may well think that I am exaggerating. But before you dismiss my viewpoint completely, look first at some examples of the dialogue between academics and practitioners having gone wrong. I will discuss three of them. They are taken from three classic areas: the CAPM, MARKOWITZ and efficient markets. I have chosen them on the basis of my own experience. In other words, I have committed all the crimes that I am about to denounce. I will perhaps be somewhat provocative at times but I will also try to be constructive.

CAPM

Let me start with the CAPM. It was a brilliant model. It changed the thinking of an industry by redefining risk. Its early proponents were responsible for the revision of the "prudent man principle", which was henceforth to be applied to the overall portfolio instead of individual securities. This change was responsible for the acceptance of the idea of American pension funds investing in high risk asset classes (private equity, emerging markets, etc.), where many other institutional investors have followed.
At the same time, the CAPM was and is also a very primitive model. It gave us three results:
- one, harmless and widely accepted, was beta as a risk measure;
- the second, first ridiculed then reluctantly accepted and now ridiculed again, was the linear relation between beta and expected return; and
- the third, an embarrassment and not even to be mentioned in a serious discussion, the notion that everybody would hold the same portfolio (which would have to be the market portfolio).

The CAPM had perhaps the misfortune to be discovered in the US at a time when international investment was not an issue and the model could be presented in a domestic context. It is therefore not surprising that it was subsequently postulated and tested separately in every imaginable country with a stock market. The results were always the same: the evidence did not support the model.

I can understand why we kept trying for a decade to find empirical support for the CAPM but I have some difficulty understanding why, 20 years later, papers are still being written about empirical tests of the CAPM for yet another market.
In this case, the problem is not with the practitioners who have rarely taken the CAPM terribly seriously but with the profession. We are apparently so desperate to have any pricing model at all that we are willing to overlook some very basic problems with the CAPM. To begin with, it is obviously not a good idea to ignore selectively some of the results of models. If we observe that one result of the model does not hold (since investors hold widely different portfolios), we should suspect that the second result (the pricing of securities in those portfolios) does not hold either. If the basic demand for securities across investors differs for whatever reason, the CAPM is unlikely to hold. This means that for the US market of the 1960s, the CAPM was a fairly reasonable model; the market was dominated by local investors holding broadly similar (domestic) portfolios. For the Swiss market of the 1990s (and even the 1960s), it is an utterly ridiculous model. The investors who are active in the Swiss market are a very heterogeneous lot (Swiss and many different foreigners) and therefore the notion that Swiss equities should be priced in a domestic CAPM context is not very reasonable.
The same problem applies to the various attempts of creating an international CAPM. If the model does not incorporate the observed home country bias of investors, its results are likely to have little normative value.
The list of the model's shortcomings is far from complete but it is already quite obvious that the CAPM as a normative model should be laid to rest. It cannot describe pricing of securities in any meaningful way. Its main purpose is as an illustration of the types of risk which are likely to be relevant for securities pricing. For that purpose, it is excellent and unsurpassed.
I believe that the unwillingness to let go of the CAPM as a serious equilibrium model is an illustration of the third factor which I mentioned. It would be so convenient if it worked that we desperately want it to work. And even some people who never believed that it worked are still quite happy to apply its results (we still use the CAPM based risk-adjusted cost-of-capital, risk-adjusted alphas and similar things).

Asset Allocation

As my second example, I have chosen MARKOWITZ' classic mean-variance optimisation procedure. He, of course, envisaged his results to be applied to individual securities in the context of "portfolio selection" as the title of his book suggests. Today, few people use MARKOWITZ'S framework for stock selection because of the vast data requirements for a realistic stock universe.
It is however widely used for asset allocation purposes, i.e., the choice of the strategic (long-term) allocation across a certain number of asset classes. Commercial software is available to make the task quite easy (e.g., Ibbotson's Encorr).
Let me describe the typical procedure which is used to determine the optimal asset allocation (I admit that there is a certain amount of caricature in the description which follows):

- First, we have to decide which asset classes to consider. The classification could use very specific asset classes (e.g. things like French mid cap pharmaceuticals) but I will stick to the usual basic ones used by Swiss pension funds (i.e., Swiss equities, foreign equities, CHF fixed income, foreign currency fixed income and CHF cash). We also have to find a way to describe our risk tolerance (the easy way out is to use the historic volatility as the target volatility but there are more sophisticated methods). While I am looking exclusively at the assets of the investor, ignoring the nature of his liabilities, the problem becomes actually worse in the context of asset-liability modelling.
- Second, we use some historical data to get estimates of the risk structure (which is normally

Table 1: Markowitz I - Historical data

Asset Class	Return (%)	Volatility (%)
CH equities	17.3	23.7
Foreign equities	11.2	22.4
CHF fixed income	5.6	4.3
FC fixed income	6.0	10.6
CHF cash	5.0	1.1

accepted as it is for forecasting). An example is shown in Table 1.

- Third, we usually generate a series of explicit return forecasts as inputs for the optimisation. The source of these forecasts could be almost anything (analysts, economists, models, etc.). They represent our best forecast of the future performance of the asset classes (Table 2 shows a set of such forecasts). With these return forecasts and the historical risk structure as inputs, we now run the optimiser.
- Next, we analyse the optimal portfolio at our chosen target risk level (Table 3 shows the results). Our reaction to these results depends on our experience. If we are new to this exercise, we are appalled by the typically unreasonable result (normally a fairly extreme corner solution). If we have done this before, we are resigned to the fact that, as usual, the optimiser needs a bit of our help.
- This help takes the form of an iterative adjustment process that continues until we arrive at a "reasonable" solution. Reasonable generally means a fairly diversified solution that also meets whatever constraints we are subject to. The adjustments are usually concentrated on the returns and, to a lesser extent, the standard deviations and correlations (see Tables 4 and 5 for one possible set of "improved" forecasts and the resulting allocations, respectively). For those without the skill or patience to perform

Table 3: Markowitz III - Optimal allocation based on forecast data

Asset Class	Weighting (%)
CH equities	4.4
Foreign equities	0.0
CHF fixed income	31.3
FC fixed income	64.3
CHF cash	0.0

Table 2: Markowitz II - Forecast data

Asset Class	Return (%)	Volatility (%)
CH equities	10.0	23.7
	17.3	*23.7*
Foreign equities	11.0	22.4
	11.2	*22.4*
CHF fixed income	5.5	4.3
	5.6	*4.3*
FC fixed income	8.5	10.6
	6.0	*10.6*
CHF cash	3.0	1.1
	5.0	*1.1*

Historical data in italics

Table 4: Markowitz IV - Improved forecast data

Asset Class	Return (%)	Volatility (%)
CH equities	10.0	23.7
	10.0	*23.7*
Foreign equities	10.2	24.0
	11.0	*22.4*
CHF fixed income	5.5	4.3
	5.5	*4.3*
FC fixed income	7.2	10.6
	8.5	*10.6*
CHF cash	3.0	1.1
	3.0	*1.1*

Original forecast in italics

Table 5: Markowitz V - Optimal allocation based on improved forecast data

Asset Class	Weighting (%)
CH equities	14.7
	4.4
Foreign equities	7.8
	0.0
CHF fixed income	52.1
	31.3
FC fixed income	25.4
	64.3
CHF cash	0.0
	0.0

Original forecast in italics

this type of fine-tuning, there is always the option of forcing the results which one desires directly by imposing a suitable set of constraints on the weights of the asset classes.

Now, we examine the long-term results by applying the lognormal model and determine various shortfall risks for a number of hurdle rates and investment horizons in which we are particularly interested (Table 6).

- Finally, we are ready to present this result to the decision makers. Naturally, we do not mention the iterative process that we used. Instead, we justify the chosen returns and the risk structure as being the result of a sophisticated blend of historical and proprietary information that provides an excellent basis for rational decision making. We offer, say, three solutions (of which only one appears reasonable). The decision-makers are impressed by the amazing technical skills that we display and also by the accuracy of our forecasts. They are more than a little bit intimidated by the incredible precision of the model especially after they are told that it won a Nobel Prize. Not surprisingly, our favourite solution is chosen with little change (which is a good thing since the solution far away from where we concentrated our improvement efforts it is as unreasonable as ever).

This is then a summary of what we call optimal asset allocation based on the MARKOWITZ optimisation process. You may think that there is nothing wrong with what I have described or that I have exaggerated. I think that I have actually been rather kind. The reality is often even worse (the use of historical returns as forecasts, etc.).

So what is wrong with this approach and does it matter? Let me first look at the question of what is wrong and get back to the question of whether it matters later. There are two types of problems: the explicit mistakes and the ignored questions. Let me start with the ignored questions:

- First, the entire exercise assumes that we know the parameters of the return distributions (means, standard deviations and covariances). In reality, we do not, all we have are some estimates of those parameters. I believe that this so-called estimation risk is one of the great neglected areas of modern finance. It is admittedly not an easy problem to deal with practically. But if there is a great deal of it, it can have a dramatic impact on the optimal allocation; thus it is particularly relevant when exotic asset classes are included in the analysis. As an illustration compare the small changes in the return forecasts in Table 4 with the dramatic changes in asset allocation which result in Table 5.
- Second, we normally have a time horizon in mind, say twenty years. In principle, we should use data of that frequency. However, since we may not have a sufficient number of twenty-year periods in our sample to estimate parameters, we automatically use one-year data. Yet a

Table 6: Markowitz VI - Shortfall risk of optimal allocation

	1 year	5 years	10 years	20 years
0% target	19.4%	2.7%	0.3%	0.0%
4% target	36.8%	22.5%	14.3%	6.5%

one-year horizon may not be terribly relevant for our purposes.

- Third, if we talk about a pension fund or other institutional investor, it is not really clear what we are optimising. Who drives the exercise? Whose risk tolerance are we talking about? That of the board of trustees, of the employer, of a retired member or of a young member? We know that they won't agree. So in which sense are we optimising?

Let us now turn to the mistakes:

- First, the iterative process is obviously nonsense. This is particularly true if we believe in our forecasts. In that case, they are an attempt to create an active allocation strategy. In such a case, we should not have been surprised to find a corner solution; after all, the optimal ex post allocation is always a corner solution. By rejecting this solution, we abandon our original forecast and the information that it supposedly represented. But if the information in our forecast is not worth using, why do we bother with it in the first place? Shouldn't we then simply use a consensus allocation?
- The typical asset allocation displays a marked home country bias. Despite the higher risk of foreign currency investment opportunities, this home country bias rarely obtains directly from the optimisation. It is generally built in artificially by tweaking the volatility structure. In reality, the home country bias is probably partially due to the fact that estimation risk may well be lower for domestic asset classes than for foreign asset classes.
- As mentioned before, even though we use a one-year horizon by default, we are generally more interested in the long-term consequences of our chosen strategy. Using a lognormal assumption, we can easily determine shortfall probabilities and the like over any horizon. But if we look at those probabilities, what do they mean (Table 6)? Looking at the probability of a negative return after twenty years, we see a comfortingly small number. But over one and five years, they look not nearly as comforting. If I concentrate on the twenty-year number, I am implicitly assuming that I will not consider revising the portfolio over the twenty years, regardless of what happens in the meantime; in other words, I will never even look at the portfolio. Is this realistic? Of course not. The problem is that we are applying a static analysis to what is really a dynamic problem. The question to which I would really like to have an answer is the following: if I pursued the optimal series of one-year strategies, what is the probability that I ever experience a shortfall over the twenty years?

MARKOWITZ optimisation is abused the way that I have described primarily because of a mixture of ignorance and wishful thinking. Occasionally, there may be an attempt to deceive but I do not think that this is the main problem. It is true of course that if we had all the necessary information, Markowitz would be the basis of a magnificent tool. As it is generally used, it is just an elaborate way to justify our preconceived notions.

Anomalies

My last example concerns anomalies, i.e., empirical results which appear to be contrary to "normal" relationships. They have had a great fascination for researchers and practitioners alike. For the first, they offer endless opportunities for empirical tests that sometimes degenerate into quite bizarre cases of data mining. For the latter, they provide an apparently simple prescription for superior performance.

For some people, their existence is an argument against market efficiency but that is really not the point. Since we have no market equilibrium model and given that any test of market efficiency is a joint test of efficiency and equilibrium model, we cannot assign blame for the anomalies to one or the other.

It is to the credit of practitioners that they have ignored most anomalies. This despite the fact that a discussion of the available profit opportunities and/or of the optimal strategy to exploit the anomaly became a standard part of the literature.

In this respect, some academics show little appreciation of the realities of the marketplace. What is the practical relevance of the weekend and turn-of-the-year effects for portfolio management? If I receive additional cash from a client on Wednesday, should I wait until Monday before I invest it? If it arrives in November, do I wait until the first Monday after Christmas to invest? Just because some t-, F- or whatever statistic is significantly different from zero does not always mean that investor behaviour is necessarily affected. Sometimes a market inefficiency can be very uninteresting from a practical viewpoint. As an example: the strong form of the efficient market hypothesis (which says that nobody makes excessive returns) can obviously be rejected. This hardly matters for the typical investor: inside information will beat the market every time, but it is hardly a good basis for a long-term investment strategy. In short, persisting anomalies may be amusing but they are mostly not very important for investors despite the academic hype that surrounds them.

On the other hand, practitioners sometimes abuse the results of these studies in order to justify a particular strategy. Dozens of American small cap managers have used my US small cap results to justify their small cap strategy. This includes managers whose definition of small cap includes companies with market caps up to USD 5 billion, well into the first quartile of the US market. It also includes managers who pursue a high turnover strategy: given the high transaction costs of small caps, the resulting costs are probably higher than any reasonable estimate of the small cap premium. Consequently, the success of these managers has more to do with stock picking capability than with the small cap premium. My results are really only relevant for a passive, buy-and-hold type strategy applied to a well-defined universe.

Conclusions

This concludes my discussion of the examples. I am not suggesting that we rip the pages containing the CAPM, Markowitz asset allocation, the SHARPE ratio and anomalies out of our textbooks. What I would like to suggest is that we reflect a bit more about the manner in which the results of the theory behind financial economics are transmitted and applied. This means that those who write the books and teach us finance abandon their missionary zeal and give us practitioners the whole story. This also means that we practitioners should take care that we make sure that when we use "modern finance" in our commercial activities, that we do not gloss over the limitations of what we are presenting.

Perhaps the best way of summarising my message is to suggest that we should strive a bit less for very precise results but strive rather more to assure understanding of the results and their consequences. If you permit me one more example: when I used to teach corporate finance, there were endless discussions about the appropriate discount rate for capital budgeting purposes, but nobody wanted to talk about the problems of forecasting cash flows. But I ask you: Have more companies gone bankrupt because they misestimated their cost of capital or because they misestimated their future cash flows. We too should try to focus on the essential.

I have not been very complementary of our profession in this article. My intentions were not to run down our profession but rather to make some suggestions for improvements for the dialogue between academics and practitioners, a dialogue from which both sides can greatly benefit.