Selected Papers . 23

Mathematical Models for Financial Management

By MERTON H. MILLER
and DANIEL ORR

GRADUATE SCHOOL OF BUSINESS
UNIVERSITY OF *CHICAGO*

MERTON H. MILLER, the Edward Eagle Brown Professor of Finance and Economics in the Graduate School of Business, is an outstanding scholar in the application of quantitative analysis to a wide variety of management problems, particularly in finance. He co-developed the celebrated Modigliani-Miller theorems on financial policy and the cost of capital. Before joining the Graduate School of Business faculty, Professor Miller had been an economist with the U.S. Treasury Department and with the Board of Governors of the Federal Reserve System, a lecturer at the London School of Economics, and a faculty member of Carnegie Institute of Technology's Graduate School of Industrial Administration. He received the A.B. degree magna cum laude from Haruard and the Ph.D. from Johns Hopkins. He recently was awarded a National Science Foundation Senior Post-Doctoral Fellowship for a year of study and research at the Center for Operations Research and Econometrics at the Universite de Louvain, Belgium.

DANIEL ORR, now on the faculty of the University of California at La Jolla, was Assistant Professor of Mathematical Economics at the Graduate School of Business of the University of Chicago until 1966. His work in the mathematical analysis of business problems has appeared in such periodicals as The Journal of Business, Economica, Quarterly Journal of Economics, and American Economic Review. He is co-author with Merton H. Miller of a series of published papers on money supply and demand, and is author of Models of the Maximizing Firm, to be published by Prentice-Hall, Inc. He received the A.B. degree from Oberlin, and the M.A. and Ph.D. from Princeton. He has been on the faculties of Princeton, Amherst, and the University of California at San Diego. He was an operations analyst for Procter and Gamble Co. and has done consulting work for Mathematica, Inc., Arthur D. Little, Inc., and The RAND Corporation.

THIS PAPER was presented at the Conference on Financial Research and Its Implications for Management, held at Stanford University June 9-11, 1966. The underlying study and report were done with the assistance of two students of the Graduate School of Business of The University of Chicago-ROBERT CRAMER, now with the United California Bank, and PETER KUYPER, now with the Richfield Oil Company.

Mathematical Models for Financial Management

ONE STREAM of current research in finance involves the extension to the field of finance of the methods and approaches that have come to be called "operations research" or "management science." Researchers working along these lines try to develop mathematical representations or "models" of typical decision-making problems in finance and, where they are given the opportunity to do so, to test and apply these models in actual decision settings. At the moment, this stream of research is still a relatively small one-really only a trickle as compared to the flood of material pouring out on the subjects of capital budgeting, or valuation. But it is a stream that can be expected to grow rapidly in the years ahead with the improvement in mathematical and computer technology and especially with the increase in the number of people who are being taught to use the tools effectively and creatively.

Rather than attempting any broad survey of work to date, this paper will present a single example of this type of research, describing both the development of the mathematical model and its application in a specific firm. Such an example can convey more graphically and more convincingly than any amount of preaching many of the important implications for management of this kind of research.

The Cash Balance as an *Inventory*

The particular financial problem involved in our example is that of managing a cash balance in conjunction with a portfolio of short term securities. And the particular mathematical model that will be used is a type of inventory model that might be called a "control-limit" model.

It may be a little startling at first to think of your firm's cash balance as just another inventory-an inventory of dollars, so to speak-but is it really so farfetched? Consider, for example, some raw material item that your company stocks and ask yourself why you keep so much of it around or why you don't simply order each day's or each hour's requirement on a hand-to-mouth basis. The answer is, of course, that this would be a very wasteful policy. The clerical and other costs involved in placing orders for the material are not trivial; and there would be further costs incurred in the form of production delays or interruptions if materials were slow in arriving or if requirements on any day should happen to be higher than had been anticipated. Why, then, not eliminate these costs once and for all by placing one big order for a mountain of the stuff? Here, of course, the answer would be that there are also costs connected with ***holding*** inventory. These would include not just the physical costs connected with the storage space and handling, but also the cost of deterioration, or of obsolescence, or of adverse price fluctuations, and especially of the earnings foregone on the capital tied up in the inventory. The inventory management problem for any physical commodity is thus one of striking a balance between these different kinds of costs; and the goal is to develop a policy in which orders will be placed on the average at just the right frequency and in just the right amounts so as to produce the smallest ***combined*** costs of ordering, of holding inventory, and of running out of stock.

Similarly with cash. If you want to add to or subtract from your inventory of cash by making a transfer to or from your portfolio of securities, there is an order cost involved, partly in the form of internal clerical and decision-making costs and partly in the form of brokerage fees, wire transfer costs and the like. In the other direction, if you try to cut down these in-and-out costs by holding large cash balances, there is a substantial holding cost in the form of the interest loss on the funds tied up in the balance. As for the costs connected with running out of cash, these are perhaps too obvious to require discussion before a group of this kind.

The Control-Limit Approach to Cash Management

Accepting the inventory analogy as valid, what form of inventory management policy would be suitable for cash balances? Here, since the typical cash balance fluctuates up and down and, in part, unpredictably, it seemed to us that the most natural approach for a wide variety of cases might be a "control-limit" policy.

How one particular kind of control-limit policy might work when applied to a cash balance is illustrated in Figure 1. We say "one particular kind" since the control-limit approach is quite flexible and many different variations can be used depending on the circumstances. The one illustrated happens to be an especially simple one and one that can be shown to be appropriate whenever the internal clerical and decision-making costs are the main costs involved in making portfolio transactions. It is also the form of policy actually used in the specific application to be described later.

The wiggly line that starts at the left at m_0 traces out the hypothetical path of a cash balance over time. As drawn, it first seems to fluctuate aimlessly until about day t_0, at which point a rising trend appears to set in. During this interval receipts are exceeding expenditures and the cash balance is building up. The buildup is allowed to continue until the day on which the cash balance first reaches or breaks through the *upper control limit* of h dollars. At this point in time-day t_1 on the graph-a portfolio purchase is made in an amount large enough to restore the cash balance to the *return point* z. Once back at z, the cash balance is allowed to wander again. No further purchases or sales are made until the balance either breaks through the upper bound at h again or until it breaks through the lower

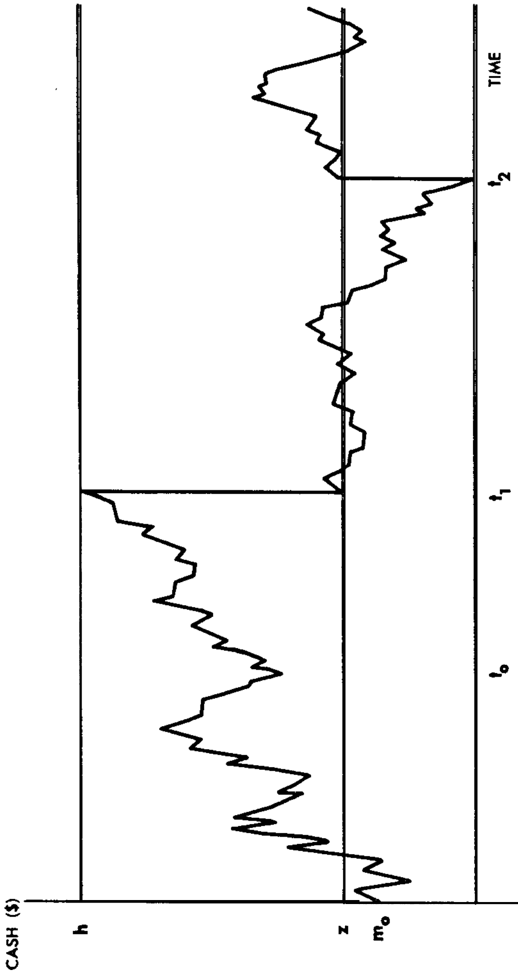

FIGURE 1

control limit, l, as at day t_2. When the lower control limit is reached, a sale of securities from the portfolio is signalled in an amount such that the balance is once again restored to the return point z.

An Optimal Solution *for a* ***Simple Special Case***

Given that this kind of control policy seems reasonable-and we would argue that it is reasonable not only in dealing with some types of cash management problems but in many other kinds of settings where there is a substantial cost in managerial intervention to restore a "wandering" system to some desired state-the task of the researcher then becomes that of applying mathematical or numerical methods to determine the ***optimal*** values of the limits. By optimal we mean values that provide the most advantageous trade-off between interest loss on idle cash and the costs involved in transfers of cash to and from the portfolio. As it turns out, there happen to be some simple, but important special cases in which these optimal values can be derived in relatively straightforward fashion and where the results can be expressed in the form of a simple, compact formula. In particular, we have been able to obtain such a formula for the optimal values of the limits for cases which meet the following conditions: (1) Where it is meaningful to talk about both the cash balance and the portfolio as if they were each single homogeneous assets.[1] (2) Where transfers between cash and the portfolio may take place at any time but only at a given "fixed" cost, i.e., a cost that is the same regardless of the amount transferred, the direction of the transfer or of the time since

1 **We have also recently been able to develop approximately optimal solutions for certain special kinds of "three-asset" models, i.e., models in which there are two kinds of securities (e.g., a line of credit and commercial paper) in addition to cash.**

the previous transfer.[2] (3) Where such transfers may be regarded as taking place instantaneously, that is, where the "lead time" involved in portfolio transfers is short enough to be ignored. (4) Where the lower limit on the cash balance is determined outside the model, presumably as the result of negotiations between the bank and the firm as to what the firm's required minimum balance is to be; and (5) Where the fluctuations in the cash balance are entirely random. There may perhaps be a trend or "drift" as it is called in this kind of analysis; but aside from this kind of simple systematic component, the day-to-day changes in the cash balance are completely unpredictable.

As for the specific formula that constitutes the solution under these assumptions, there is little point in discussing it any length here. The complete derivations and other details can be found in a recently-published article.[3] It might, perhaps, just be worth noting here that the solution defines the limits in terms of the fixed transfer cost, the daily rate of interest on the portfolio and the variability of daily changes in the cash balance (exclusive of changes related to the portfolio). As would be expected, the higher the transfer cost and greater the variability the wider the spread between the upper and the lower limits; and the higher the rate of interest, the lower the spread. There are, however, some surprises. In particular, for the "no-drift" case, it turns out that despite the fact that the cash balance

2 Simple solutions also have been developed for the case in which the cost is not fixed but proportional to the amount transferred. More complicated, mixed cases involving both a fixed and a proportional component have been analyzed by our colleagues G. Eppen and E. Fama who have developed a very flexible method of obtaining numerical values for the limits under a wide variety of circumstances.

3 "Model of the Demand for Money by Firms," by M. H. Miller and D. Orr, Quarterly Journal of Economics (forthcoming).

is equally likely to go up or down, and it's equally costly to buy or to sell securities, the optimal return point z-the point at which the average long-run costs of operating the system are lowest-does not lie midway between the upper and lower limits. Instead, it lies substantially below the midpoint. To be precise, it lies at one-third of the way between the lower and upper bounds and it stays at the one-third point regardless of the numerical values that are assigned to the transfer costs or to the daily rate of interest that can be earned on the portfolio. As these values are changed, the whole system expands or contracts, but the relation between the parts remains the same.

A Test Application of the Simple Model

Your initial reaction is likely to be that this model and the assumptions on which it was based are much too special and restrictive to have any important applicability to real-world problems. In management science, however, as in science generally, it is rash to pass judgment on the range of applicability of a model solely on the basis of the assumptions that underlie it. Mathematical models often turn out to be surprisingly "robust" and insensitive to errors in the assumptions. The only safe way to determine how well or how poorly a model works is to try it out and see.

In obtaining the basis for this kind of test of the model we were extremely fortunate in having the active collaboration of Mr. D. B. Romans, Assistant Treasurer of the Union Tank Car Company.[4] Mr. Romans had seen an earlier version of our original paper and was struck by the similarity between the model and his own policies in putting his firm's idle

4 **We have also benefited greatly from discussions of cash management problems and practices with several officers of the Harris Trust and Savings Bank of Chicago. We hope that they will benefit too from this chance to see how the problem looks from the other side of the account.**

cash to work. The systematic investment of idle cash in short term, money market securities was a relatively new program for his company-one that he had instituted only about a year previously. The interest earnings for that year were quite large, not only in relation to the costs involved but to the total budget of the treasurer's department. Now that the year's experience had been accumulated, he wanted to go back over the record, to study it in detail and to see whether any changes in practice might be suggested that would make the operation even more profitable. He felt, and we agreed, that the model might be extremely helpful in this kind of evaluation. If the model did seem to behave sensibly when applied to the company's past cash flow then it might be used to provide an objective standard or "bogey" against which past performance could be measured.

Since mathematical modeling of business decisions is still quite new, and since few people outside the production area have had much direct connection with it, it is perhaps worth emphasizing that at no time was it intended or contemplated that a model should be developed to do the actual on-line decision-making. The purpose of the study was to be *evaluation* by the treasurer of his own operation. This is a valuable but unglamorous use of models that tends to be overlooked amidst all the hoopla of the Sunday supplement variety surrounding the subject of automated management. An important point that must be kept in mind about mathematical models is that they are not intended to *replace* management-though like any other technological improvement they sometimes have that effect-but that they provide managers with new tools or techniques to be used *in conjunction with* other managerial techniques (including good judgment) for improving over-ah performance.

The Setting of the Operation

Since our objective was to compare the model's decisions over some trial period with those of the Assistant Treasurer, the first step was to examine carefully the setting in which he actually operated and to see how closely or how poorly the circumstances matched the assumptions of the model. As would be expected, the results were mixed. On the one hand, there were some respects in which the assumptions fit quite well. The Assistant Treasurer did behave, for example, as if he were in fact controlling only a single-central cash balance. Note the phrase "as if," because as a matter of fact the firm does have many separate balances in many banks. For purposes of cash management, however, the Assistant Treasurer works with one single balance representing the free funds that he can marshal throughout the system without regard to the particular banks they happen to be in at the moment (or where the funds derived from a portfolio liquidation must ultimately be routed).

It was also clear that there were substantial "order costs" involved in making portfolio transfers. In the case of a portfolio purchase, for example, some of the main cost components include: (a) making two or more long-distance phone calls plus fifteen minutes to a half-hour of the Assistant Treasurer's time, (b) typing up and carefully checking an authorization letter with four copies, (c) carrying the original of the letter to be signed by the Treasurer, and (d) carrying the copies to the controller's office where special accounts are opened, the entries are posted and further checks of the arithmetic made. It is hard to establish a precise dollar figure for these costs, but at least the approximate order of magnitude for a complete round trip is probably somewhere between \$20 and \$50. That this is not a trivial amount of money in the present context becomes clear when you remember that interest earnings at the then prevailing level of interest rates were running at about $10 per day per hundred thousand dollars in the portfolio and that his average size of portfolio purchase during the test period was about $400,000.

Not surprisingly, we also found that there was a considerable amount of randomness or unpredictability in the daily cash flow. In fact, the Assistant Treasurer did not even attempt to forecast or project flows more than a day or two ahead except for certain large recurring outflows such as tax payments, dividend payments, sinking fund deposits, transfers to subsidiaries and the like; and even here the forecasts were made more with a view to deciding the appropriate maturities to hold in the portfolio than as part of the cash balance control per se. As for the "drift" or trend, analysis of the cash flow over the 9-month test period showed no evidence of any significant drift in either direction.

As opposed to these similarities between the assumptions of the model and the reality of the firm's operation there were very definitely a number of respects in which the fit was much less comfortable. The model assumes, for example, that when the lower bound on cash is hit or breached, there will be an immediate sale of securities out of the portfolio to make up the cash deficiency. The Assistant Treasurer, however, followed a policy of buying only non-marketable securities and holding them to maturity primarily because he wanted to try his new cash management program without requiring any change in the company's standard accounting procedures. Hence, if a large net cash drain occurred unexpectedly on a day on which he had no maturing security he simply let his cash balance drop below his normal minimum which he and his banks regarded as an "average minimum" rather than as the strict minimum contemplated by the model.

A discrepancy between the model and reality that was more disturbing appeared when we constructed the frequency distribution of daily cash changes by size of change over the g-month sample period of 189 working days. The distribution of these daily changes is shown in graphic form in Figure 2. The logic of the model requires that this distribution be at least approximately of a form that statisticians refer to as "normal" or "Gaussian." A hasty glance at the figure might lead one to conclude that this requirement is met. Closer study reveals, however, that not only is the distribution not normal, but it almost

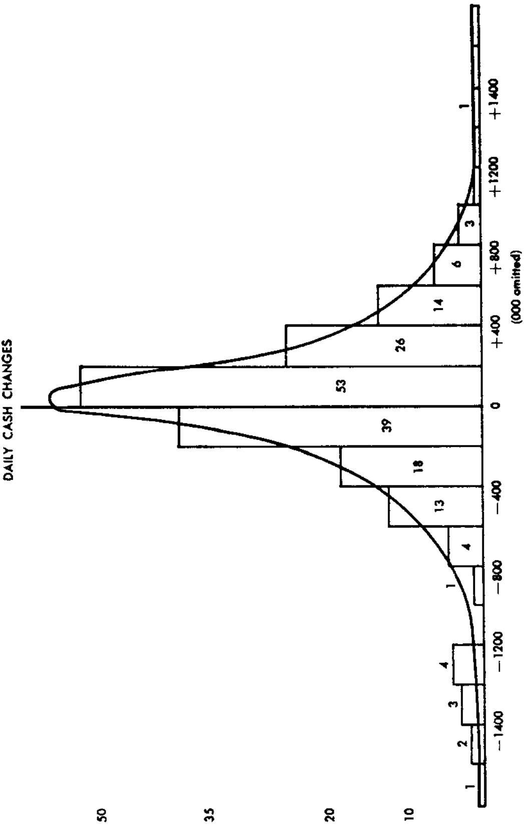

FIGURE 2

seems to be a member of a particularly ill-behaved class of "fat-tailed" distributions that have come to be called "Paretian distributions."[5] In these distributions-which may be familiar to those of you who have been following the debate about random walks in the stock markets-large changes occur much more frequently than in the case of the normal distribution. In fact, the frequency of large changes is so much greater that we were quite uncertain as to whether the model would behave in even roughly sensible fashion or whether it would simply find itself being whipsawed to death by the violent swings through the control range. As indicated earlier, however, there is only one way to tell; and that's by trying it out and seeing what happens.

The Test of the Model Against the Data

To get a close basis for comparison with the Assistant Treasurer's actual decisions it was decided to run the model under various alternative assumptions about the true value of the transfer cost. That is, we would start with a conservatively high value of say $90 per transfer; compute the optimal upper limit h and return point z; run the model against the actual data and tabulate its portfolio purchases and sales. If, as expected, the model made fewer transfers than the Assistant Treasurer, then we would go back; use a lower value for the costs; recompute the new optimal

5 We say "almost seems to be" because despite the conspicuously fat tails, the distributions as computed cumulatively month by month remain roughly similar with no tendency for the tails to get fatter and fatter over time as in a true Paretian process. The Paretian-like tails are mainly the reflection of such large, but relatively controllable and definitely size-limited items as dividends, taxes, transfers to and from subsidiaries and the like.

6 *Random Walks in Stock-Market Prices,* Selected Papers No. 16, Graduate School of Business, University of Chicago.

limits and so on until we had finally forced the model to make approximately the same number of transfers over the sample interval as the Assistant Treasurer himself. Then, assuming the model was behaving sensibly, we could compare and contrast their patterns of portfolio decisions over the interval as well as get at least some rough idea of what figure for the cost of a transfer the Assistant Treasurer was implicitly using in his own operation.

The only difficulty encountered in implementing this straightforward kind of test was in the matter of deciding precisely how many transfers the Assistant Treasurer should be regarded as having made. Because of his policy of holding only non-marketable issues, his portfolio tended to be of quite short average duration. Hence there were inevitably days on which he had a maturity that proved to be too early. If he had merely rolled these issues over there would have been no problem; we would simply have washed that transaction out and not counted either the maturity or the reinvestment as a transfer. But it is clearly not always efficient just to roll over the maturing issue. Given that a purchase must be made anyway on that day, it would be wise to pick up any additional cash that also happened to be lying around, even if the amount involved would not have been large enough by itself to have justified incurring a transfer cost. Accordingly, we decided not to count any transfers on rollover days unless the Assistant Treaurer indicated that the balance was so large even without the maturing issue that he would almost certainly have bought anyway (in which case he would be charged with the purchase, but not the sale). Similarly with the case of net sale days. If there was a larger maturity on a given day than was actually needed to meet the cash drain and if some small part of the excess proceeds were rolled over, then he was charged with a sale, but not a purchase. By this criterion we were able to agree on a figure of 112 total transactions by the Assistant Treasurer during the 189 test days of which 58 were purchases and 54 were sales (maturities).

The Results **of** *the Test*

When we commenced the trial-and-error-process of matching the total number of transactions by the model with those of the Assistant Treasurer our hope was that the model might be able to achieve an average daily cash balance no more than say 20 to 30 per cent above the Assistant Treasurer's average. We felt that if we could get that close and if the model did behave sensibly, then there was a very real prospect of being able to use the model as a bogey against which to measure and evaluate actual performance. As it turned out, however, we found that at 112 transactions, the model not only came close but actually did better-producing an average daily cash balance about 40 per cent lower than that of the Assistant Treasurer ($160,000 for the model as compared with about $275,-000). Or, looking at it from the other side, if we matched the average daily cash holdings at $275,000, the model was able to reach this level with only about 80 transactions or about one-third less than the 112 actually required.

It can be argued, of course, that this sort of comparison is unfairly loaded in favor of the model, not only because it was applied on a hindsight basis, but because the transfer costs would actually have been higher for the model than the simple matching of total numbers of transfers would seem to suggest. The Assistant Treasurer, it will be recalled, never really sold a security; he merely let it run off. Hence the model would have had to incur additional costs on at least those sales that did not occur on the easily forecastable, large-outflow days. Check of the numbers involved showed, however, that the model would still have dominated in terms of net interest minus transfer costs over the sample period even if these extra costs of liquidation were included on every sale. And, of course, that is much too extreme an adjustment. Many of the actual sale days of the model coincided with the large out-flow days and appropriately maturing securities could have been purchased to hit these dates. In fact, the post-mortem showed that about half the model's sales took place on days when the Assistant Treasurer also "sold" and nearly 80 per cent occurred either on the same day or within one day either way of a day on which he scheduled a maturity.

Furthermore, the model too is operating under some handicaps in the comparison. At no time, for example, did the model ever violate the minimum cash balance marked on the Assistant Treasurer's work-sheets, whereas no less than 10 per cent of his total dollar days invested were represented by the cash deficiencies on the days in which he let his balance dip temporarily below the minimum. In addition, the model did not receive instructions to change its policies before weekends and holidays. The Assistant Treasurer, on the other hand, always knew when it was Friday and was thus able to sock away additional amounts on which he could get two extra days' interest.

All in all then the comparison would seem to be basically a fair one; and it is a tribute to the Assistant Treasurer's personal and professional character that he never became ego-involved in the comparison or wasted time alibiing. He was concerned about one thing and one thing only: how to do an even better job.

The Comparison **of** *Operating Policies*

With this question in mind, we then went on to make a detailed comparison of the actual decisions with those called for by the model. The complete record of these comparisons is, of course, too long and too specialized to be spelled out at length here, but there are at least a few simple contrasts that can be presented to illustrate the sorts of things that turned up.

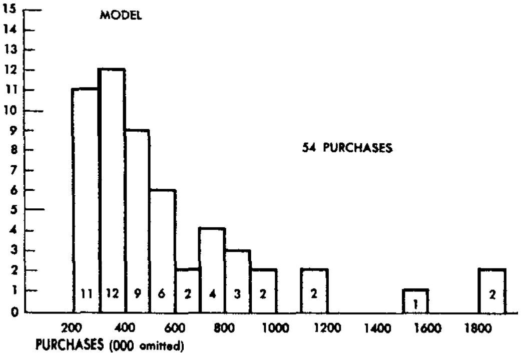

FIGURE 2a

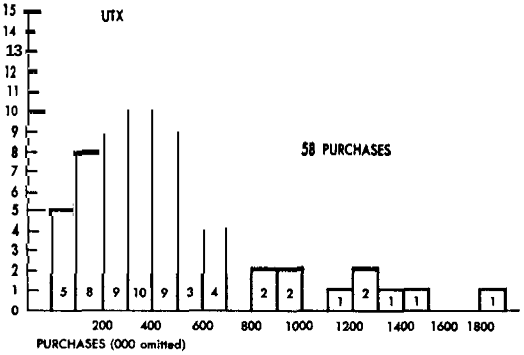

FIGURE 2b

Figures 2a and 2b, for example, show the frequency distributions of portfolio purchases by size of purchase for the model and for the Assistant Treasurer. Notice that even though we have forced the total number of transfers to match, the model makes somewhat fewer purchases (54 as against 58) and does so in considerably larger average size (about $600,-000 as compared with only $440,000). The difference in operating policy is particularly striking at the lower end of the size scale because of the rigid rule built into the model that keeps it from ever buying in units smaller than h-z, which was about $250,000 when the model was set to produce 112 transfers. The Assistant Treasurer, by contrast, made about 13 purchases (or nearly 25 per cent of his total purchases) in amounts smaller than that size including 5 in amounts of $100,000 or less. Even allowing for the fact that some of these small transactions were for weekends, the total impression conveyed is one of an excessive amount of small-lot purchasing activity. This impression was further reinforced both by the very low implicit transfer cost that was necessary to force the model to make 112 transfers as well as by the fact that more than 90 per cent of the total interest earnings achieved by the model with 112 transfers could have been attained with only about 50 total transfers. Of these 50, moreover, only some 20 were purchases and all were of fairly large size.

Even more revealing are Figures 3a and 3b which show the distribution of the closing cash balance by size on days when no portfolio action was taken in either direction. Notice again that because of its rigid upper limit the model never lets the cash balance go above h which in this case is about $400,000. The Assistant Treasurer, however, seems to be much less consistent in this respect, having foregone no less than 23 buying opportunities of this amount or larger including 3 of over a million dollars. When and why so many opportunities were missed is still not entirely clear. Part of the trouble undoubtedly stems from the fact that the Assistant Treasurer has many other responsibilities and cannot always count on being at his desk at the time of day when the decision has to be made. And without actually interrupting to construct his worksheet, there is no way for him

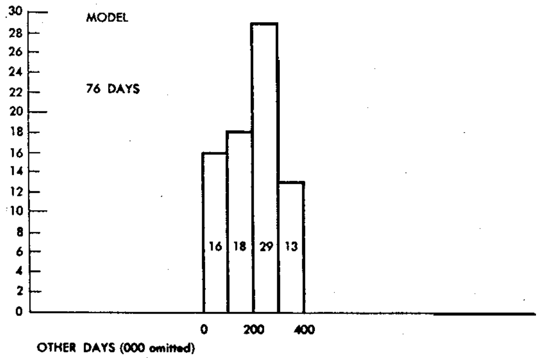

FIGURE 3a

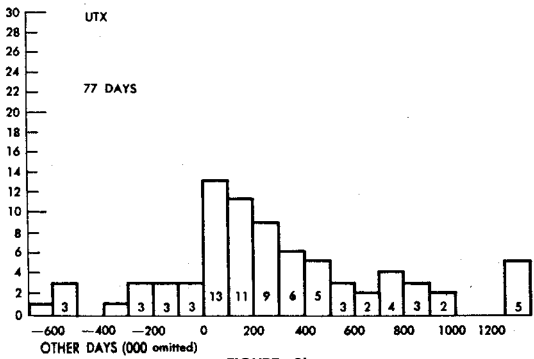

FIGURE 3b

to determine whether an interruption of his other work would really be profitable. Hopefully, however, by making his limits more explicit (in the spirit of the model) and by delegating to others much of the purely mechanical task of monitoring these limits, he will be able to achieve in the future a significant reduction in the size and frequency of these lost opportunities.

Conclusion

We have tried here to present a concrete example of how mathematical methods can be and are being applied to management problems in the field of finance. The example happens to be a particularly simple one. But it does at least serve to illustrate very neatly a number of points about this kind of research that senior financial managers would do well to keep in mind.

First, it is important for financial managers to disabuse themselves of the notion that there is something special or unique about financial problems. In particular, we have seen that what is commonly regarded as a peculiarly financial problem-to wit, managing the cash balance and a portfolio of liquid securities-turns out to be nothing more than an inventory problem.

Second, mathematical models of decision or control problems should not be thought of as something fundamentally different from ordinary management principles or techniques. They are merely more disciplined and systematic ways of exploiting these principles. In particular, control-limit models of the kind we have seen here-and remember that many additional variations are possible-are essentially extensions of the fundamental notion of "management by exception."

Third, be careful not to prejudge mathematical models solely on the basis of the lack of literal realism in the assumptions underlying them. To develop a workable model, simplifications-sometimes, extreme simplifications-must be made. But, if it has been properly conceived, a simple model may still perform extremely well. It is not a matter of getting something for nothing; rather that the gains made by doing a good job on the really essential parts of the problem are often more than large enough to offset the errors introduced by the simplifications (errors, incidentally, that often cancel out).

Finally, remember that there is a trade-off between improving decision-procedures and improving the information and forecasts used in arriving at the decisions. In the present instance, for example, we saw a case in which a model that assumed the cash flow to be completely random was still able to do a very successful job of decision-making. Nor is this result unique or exceptional. The slogan everywhere today is "more, better and faster information for management." We suspect, however, that thanks to the computer, many firms may already be in the position of having more, better and faster information than they can use effectively with present management techniques. There is likely to be as much or more real pay-off in the years ahead in rationalizing and improving decision procedures than there is in simply trying to get an even bigger bang from the information explosion.