GAINS FROM REAL-TIME FLOOD CONTROL OPERATION

Stuart S. Schwartz
Water Resource Systems Engineer

Report 88-2

April 1988

Interstate Commission on the Potomac River Basin
6110 Executive Boulevard, Suite 300
Rockville, Maryland 20852-3903

Table of Contents

List of Figures

List of Tables

Acknowledgments

Many individuals contributed both directly and indirectly to the completion of this work. Dr. Roland C. Steiner sustained the support and flexibility that allowed this work to be pursued. Support for portions of this work were provided by the Maryland Department of Natural Resources, Water Resources Administration, and the United States Geological Survey. Dr. J. Hugh Ellis generously provided time on the Cornell National Science Foundation Super computer. Dr. J. Hugh Ellis, Dr. James A. Smith and Dr. Roland C. Steiner offered thoughtful comments and careful review of an earlier version of this manuscript. The assistance and support that was received in completing this work is gratefully acknowledged. Responsibility for all assertions, errors, and conclusions contained herein lies solely with the author.

This report resulted from the enthusiastic support and steadfast encouragement of original and creative work that was fostered at ICPRB by Paul W. Eastman. Those who had the privilege of knowing him are richer for the experience.

This report has been prepared by the Interstate Commission on the Potomac River Basin. Funds for this report were provided by the United States Government, the U.S. Environmental Protection Agency, and the signatory bodies of the ICPRB: Maryland, Pennsylvania, Virginia, West Virginia, and the District of Columbia. The opinions expressed are those of the authors and should not be construed as representing the opinions or policies of the United States or any of its agencies, the several states, or of the Commissioners of the Interstate Commission on the Potomac River Basin.

Executive Summary

The Potomac River is the second largest source of freshwater to the Chesapeake Bay. Additional reservoir storage on the Potomac River could provide a supplemental source of freshwater for the Chesapeake Bay during low flow periods in the summer and late fall. If flood control storage on the Potomac could be operated more efficiently, reservoir storage in existing impoundments could be operated to supplement freshwater flows without sacrificing any flood protection.

This study examined the potential for improving the operation of flood control reservoirs through the combined use of stochastic real-time reservoir operating rules and real-time forecasting. Where operating gains can be realized in flood control operation, the reliability of flood protection can be increased, flood damages can be reduced, and the volume of storage required to provide flood protection can be reduced, freeing storage in existing reservoirs for alternate uses.

- Reservoir releases must hedge against uncertain hydrologic conditions during a flood. Operating rules that incorporate the accuracy of the runoff forecast in release decisions will deliver increased protection from an improvement in forecast accuracy.

- The operating forecast represents a probability density function of runoff. Operating rules that explicitly use this density function in determining releases reduce the cost of risk taking through non-structural hedging.

- Releases that explicitly balance operating risks in real-time can achieve true multipurpose reservoir operation through the operation of jointly allocated reservoir storage.

This study considered the reallocation of 5,600, 10,300, and 12,200 acre-feet (AF) of flood control storage in Jennings Randolph Reservoir to conservation purposes. Four levels of forecast skill were simulated and used to compare the operating consequences of both a fixed allocation of flood control storage, as well as a joint allocation in which a portion of the flood pool was allocated to both flood control and conservation purposes.

The results indicate that great potential exists for utilizing real-time forecasting and operating procedures to support additional uses of existing flood control storage. Existing cost-effective technology can provide significant skill in estimating the probability density function of runoff during a flood. Real-time operating rules that can use probabilistic forecast information are a powerful tool for improving the utilization of existing reservoir storage.

On the Potomac river up to 12,000 AF of reservoir storage could be made available to enhance the flow of freshwater to the Chesapeake Bay without compromising existing reservoir operation. The following conclusions can be drawn:

- The reallocation of a portion of the existing flood control pool in Jennings Randolph reservoir to conservation purposes will increase the downstream flood damages. The use of improved forecasts, in conjunction with operating rules that can respond to these forecast improvements can reduce, and in some cases eliminate the loss of flood protection resulting from a storage reallocation.

- For every storage reallocation considered in this study, operating with forecasts that are updated during the flood event significantly reduces flood damages compared to operating strategies with constant hedging.

- The forecast improvements that yield the greatest marginal reductions in flood damages, are also the improvements that are most easily achieved. Additional incremental improvements in forecast skill will be increasingly difficult to achieve, and will have small but positive incremental value.

- Operation of jointly allocated storage significantly improves the benefits that can be achieved through the use of stochastic real-time operating rules and real-time forecasting. For the operating simulations considered in this report, maximum forecast skill combined with the joint allocation of 12,200 AF of flood control storage to jointly allocated storage could reduce the Luke flood peak by over 6,000 cfs for the Bloomington Standard Project Flood while providing an additional 12,200 AF of conservation storage to enhance the flow of freshwater to the Chesapeake Bay.

- Operation with stochastic real-time operating rules and improving real-time forecasts dramatically increases the flood protection offered per acre-foot for all storage allocations. The increased productivity of flood control storage achieved with real-time operation and forecasting can more than offset the loss of flood protection resulting from a reallocation operated with constant hedging.

1. INTRODUCTION

This report examines real-time flood control operating rules that allow existing reservoir storage to be used to enhance freshwater inflow to the Chesapeake Bay during low flow periods in the late summer and fall. One way to provide additional low flow augmentation benefits is to reallocate a portion of the existing flood control storage in a multipurpose reservoir to conservation purposes. Such reallocation decisions are usually analyzed by comparing the loss of flood protection to the increased benefits from low flow augmentation resulting from a marginal reallocation of flood control storage. The operating rules used to guide release decisions will determine both the effectiveness of flood control operation as well as the cost of a marginal reallocation of flood control storage. Operating rules that are able to use flood storage more efficiently will provide a higher level of overall benefits. Improvements to flood control operating rules could support the reallocation of reservoir storage to conservation purposes while preventing an overall loss of flood protection.

This report considers flood control operating rules that prescribe stochastically optimal releases in real-time. Real-time flood control operating decisions are made based on probabilistic real-time forecasts of reservoir inflow and runoff. In this stochastic environment, determining the optimal reservoir release in real-time is equivalent to finding the release that is optimal with respect to hedging against hydrologic uncertainty.

Hedging decisions forego current benefits with certainty in order to reduce the magnitude and probability of future costs. In flood control operation the cost of hedging is manifested as increases in both storage utilization and flood damage, compared to operating results that could be realized if the flood hydrographs were known with certainty. Improved forecasting in conjunction with stochastic operating rules can reduce the costs of hedging. Optimal hedging rules using improved forecast information make more favorable tradeoffs of storage utilization and flood damage feasible. The combination of forecast improvements and stochastic operating rules can reduce the flood damage cost of reallocating flood control storage.

Section 2 provides an overview of multipurpose reservoir operation. Most multipurpose reservoirs are more accurately described as multipurpose impoundments, behind which a number of single purpose reservoirs are created. Although the capacity of the impoundment is a binding constraint linking the allocations in a multipurpose project, the benefits available from each authorized project purpose in a multipurpose impoundment are usually analyzed independently. Independent analysis of each storage allocation leads to an incremental volume of storage being provided in the operating pool for each authorized project purpose, in order to reliably provide the intended operating benefits. The combination of operating rules and the allocation of storage determines the reliability with which benefits will be provided. Operating rules that use real-time forecasts to reduce the risk from hydrologic uncertainty allow reliable operation to be provided using less reservoir storage.

Section 3 considers flood control operations that minimize the flood stage at a downstream damage center. Deterministic hydrology is used to identify the maximum level of protection that could be achieved. The cost of hydrologic uncertainty is quantified by comparing stochastic optimal releases to the optimal deterministic regulation. The differences between the deterministic and stochastic optimal releases demonstrate the consequences of hedging when operations must be based on probabilistic forecasts of reservoir inflow and runoff. Three sources of hedging behavior are identified from this comparison, and the relative dominance of each is identified as conditions change throughout the operating horizon.

Section 4 considers the role of hydrologic uncertainty in real-time flood control operation. In flood control operation the important consideration in forecasting is not the mean performance of the forecast, but the rate at which forecast uncertainty changes throughout the operating horizon. Two metrics are proposed to quantify the change in hydrologic uncertainty during a flood event. The three sources of uncertainty identified in the operating comparison of section 3 are discussed in greater detail. The changes in the probability distribution of the forecast are quantified by the rate of change in modal probability and the rate of change in forecast entropy. Two simple examples are developed to suggest the way in which hydrologic uncertainty will decrease during a flood.

Section 5 simulates the change in forecast uncertainty during a flood event to demonstrate the changes in real-time flood control operation resulting from improved real-time forecasting. Real-time flood control operation of Jennings Randolph reservoir is simulated using the current allocation of storage. Reallocations of a portion of the existing flood pool are then simulated in real-time, and compared to reallocated operation with forecast improvements. The equivalence of structural storage and improved forecast skill is demonstrated using real-time stochastic operating rules.

Section 6 considers the joint reallocation of reservoir storage to both flood control and conservation purposes. Hedging in flood control operation is extended to hedging in the operation of a joint allocation of storage. The applicability of the stochastic flood control rules to the operation of jointly allocated storage demonstrates the robustness of the real-time operating problem. Successful operation of jointly allocated storage can increase the benefits available for both authorized project purposes through non-structural hedging. Using real-time forecast information, the stochastic operating rules realize non-structural hedging by balancing the operating risks in each release decision. Applying the stochastic operating rules to jointly allocated storage achieves true multipurpose operation of a multipurpose reservoir.

2. Multipurpose Reservoirs

Most multipurpose reservoirs are more accurately described as multipurpose impoundments, behind which a number of single purpose reservoirs are created. Both the design and operation of multipurpose reservoirs maintain separate allocations of storage for each authorized project purpose. A proportional fraction of reservoir inflow is credited to each authorized storage allocation, reinforcing the independent, single purpose nature of the allocations in a multipurpose project. Operation of each allocation is largely independent of all others, using separately identified single purpose operating rules. The reliability with which benefits can be provided from a multipurpose reservoir is determined by the combination of the storage volume allocated for each authorized project purpose, and the operating rules that are used to determine releases from each allocation. Operating rules that use real-time forecasts to reduce the risk from hydrologic uncertainty allow reliable operation to be achieved using less reservoir storage.

Reliable operation is achieved by providing a safety factor in each storage allocation. This safety factor takes the form of an incremental volume of storage that will rarely be used. Providing this storage gives the reservoir operator the means to hedge against hydrologic uncertainty without compromising system benefits. For example when water supply conditions are not known with certainty, hedging storage provides the operator with a reserve of water that can be released in excess of expected demands in order to avoid a shortfall. In flood control operation runoff with marginal damage potential can be stored in hedging storage to provide reliable flood protection.

In this report reservoir storage used to protect against hydrologic uncertainty is referred to as hedging storage. Hedging storage is not explicitly identified as a distinct storage volume in sizing a reservoir. Rather, the use of an incremental volume of storage is implicit in the operating rules used to determine the required volume of storage. The use of safety factors in reservoir release rules equates reservoir reliability with reservoir storage through the use of hedging storage. The reliability or performance criterion for the system can be thought of as the probability that sufficient storage will be available to support operations that realize design benefits. Structural hedging in multipurpose reservoirs achieves the design reliability by incorporating infrequently used volumes of storage in each allocation.

2.1 THE ROLE OF OPERATING RULES

Reservoir operating rules will determine both the required volume of storage in design studies, as well as system performance in operating studies. Most operating rules for both flood control and conservation purposes can be characterized as target-seeking rules. A target discharge is identified, and reservoir releases are adjusted to meet the target plus a safety factor. (For flood control operation the target is a maximum stage or discharge, while conservation targets will impose a lower bound on releases). The capacity and current contents of the reservoir are not explicitly considered in determining the current release unless storage levels approach extreme levels (i.e nearly full for flood control or nearly empty for conservation storage). Forecast information may be used to provide an expected value for the calculation of the target, but the full probabilistic information contained in the forecast is usually not employed. The risk embodied in the use of a probabilistic forecast is assumed to be mitigated by the use of a safety factor. Hedging against the probabilistic runoff forecast is achieved through the use of a safety factor.

A target seeking operating rule determines the release from a reservoir as the sum of an operating target plus a safety factor:

$$R(t) = R_T(t) + R_S(t) \qquad (2.1)$$

where
$R(t)$ is the reservoir release at time t
$R_T(t)$ is the operating target for time t
$R_S(t)$ is the safety factor for time t.

The use of a safety factor in determining reservoir releases provides benefits that exceed the nominal performance criteria under normal conditions. This provides a cushion or safety factor so that a degree of resiliency is achieved allowing the performance criteria to be satisfied under unexpectedly adverse conditions.

Traditional approaches to sizing reservoirs for both flood control and conservation purposes use critical period analysis. In critical period analysis reservoir operations are simulated to determine the volume of storage required to satisfy a design performance criteria. The required volume of storage satisfies the performance criteria for a critical hydrologic sequence taken from an historical or synthetic record. The return period associated with the critical hydrology establishes an implicit level of reliability for the reservoir. The use of target-seeking rules in reservoir sizing studies effectively builds reliability, referred to as hedging storage in this report, into a storage allocation. The use of independent target seeking rules to size each allocation within a multipurpose reservoir implicitly includes hedging storage in the allocation for each authorized project purpose.

Storage design with target-seeking rules

Critical period analysis is used to size reservoir storage by simulating operation over a design hydrologic event. For conservation purposes simulated withdrawals are made from an infinite reservoir to supplement natural inflow. The net cumulative withdrawals establish the minimum volume of storage required for successful operation during the design hydrology. Flood control storage is sized by simulating operation over selected flood hydrographs. The maximum storage used in the operating simulations is the minimum volume of storage needed to provide flood protection for the design event.

For conservation purposes the required volume of storage is

$$S = \sum_{t \in T_C} (R(t)-I(t)) \qquad (2.2)$$

where $I(t)$ is the reservoir inflow during time t, and T_C is the set of time periods constituting the critical period.

The minimum required flood control storage is

$$S = \sum_{t \in T_C} (I(t)-R(t)) \qquad (2.3)$$

From (2.1)-(2.3) it is clear that the storage allocation determined by critical period analysis can be divided into an allocation due to the target release (target storage) and an allocation due to the safety factor (hedging storage). For conservation purposes the target storage is

$$S_T = \sum_{t \in T_C} (R_T(t)-I(t)).$$

For flood control storage the target storage is

$$S_T = \sum_{t \in T_C} (I(t)-R_T(t) \quad .$$

The hedging storage is

$$S_H = K \sum_{t \in T_C} R_S(t)$$

where K = 1 for conservation purposes and -1 for flood control, reflecting the negative sign of the flood control safety factor.

The class of the target seeking release rule is general enough to include many operating strategies used in practice. Operating rules will differ in the way in which target releases and safety factors are determined. Both target and safety factors may be determined over a broad range of time scales, ranging from seasonal or annual to hourly operation. Targets and safety factors may be constant average quantities or the product of large complex simulation or optimization models. Despite the variety of techniques and formulations that have been used to determine reservoir releases for many purposes, most operating rules can be written as target seeking rules as in (2.1).

Target seeking rules used to independently size storage in multipurpose impoundments will provide hedging storage within each allocation. Real-time forecasting combined with real-time operating rules can improve multipurpose reservoir operation in three ways. Forecasts with increased accuracy reduce the magnitude of the safety factor and the required volume of hedging storage. Real-time operating rules that hedge against uncertainty in a non-structural way can provide the equivalent level of reliability using less hedging storage than target-seeking rules. For flood control and conservation purposes in which the need for storage occurs under different hydrologic regimes, the simultaneous need for hedging storage for both purposes is extremely unlikely. Real-time rules that balance operating risks can share the combined allocations of hedging storage, achieving true multipurpose operation of a single volume of storage.

2.2 NON-STRUCTURAL HEDGING

In this report stochastic operating rules are developed for real-time operation of flood control storage. Stochastic operating rules explicitly use the probabilistic information available in the current hydrometeorologic forecast, as well as the current state of the reservoir-channel system, to determine the current period's release. The operating rules balance hydrologic risks using the probability distribution of runoff as the forecast. The active use of forecast information to balance risks in real-time operation is referred to as non-structural hedging.

Structural hedging leads to the effective allocation of incremental reservoir storage in order to achieve design reliability under target oriented operating rules. With operating rules that do not actively utilize the probabilistic description of the hydrometeorologic system in determining releases, constant hedging stores more flood waters, increasing the storage required to meet the design level of reliability. The operating consequences of forecast uncertainty are managed with reservoir storage.

Non-structural hedging uses stochastic operating rules to take advantage of probabilistic forecasts, making non-damaging releases whenever possible. Hydrologic variability is shifted between short-term channel storage (through reservoir releases) and reservoir storage as hydrologic conditions change. The releases prescribed by real-time stochastic operating rules are guided by forecast distributions of reservoir inflow and runoff. Non-structural hedging uses the probabilistic forecast to release inflow when feasible, reducing the reservoir storage needed to provide flood damage reduction. Using stochastic operating rules the system reliability associated with hedging can be provided through non-structural means; hedging storage can then be used for other purposes without compromising flood control reliability.

Non-structural hedging using real-time forecasts and real-time stochastic operating rules can improve the efficiency of storage utilization in multipurpose reservoirs in two ways. First, stochastic operating rules using real-time forecasts can achieve reliable flood damage reduction using less storage. The incorporation of hedging storage in each allocation suggests similar gains can be realized from operations using improved forecasts in conservation purposes as well. For water supply, system operating rules allow reservoir storage requirements to be directly reduced in response to improvements in forecast skill (Schwartz and Sheer, 1984). Second, the effective allocation of hedging storage to each allocation understates the reliability of the system. The probability of simultaneously needing the hedging storage in both conservation and flood control allocations is extremely small. The hedging storage contained in conservation allocations during flooding conditions is effectively providing reliability against an event of extraordinary return period. By balancing the operating risks represented with real-time probabilistic forecast information, the same volume of storage can be operated for several purposes. Non-structural hedging uses real-time forecast information to achieve true multipurpose operation.

The remainder of this report considers the value of improved real-time forecasts and real-time operating rules in flood control operation. The need to hedge as well as the cost of hedging is quantified in the following section.

3. FLOOD CONTROL OPERATION

In this section the general characteristics of flood control operation are considered. Optimal deterministic operation of Jennings Randolph Reservoir is simulated for the Bloomington Standard Project Flood (SPF). Real-time operation for the Bloomington SPF is simulated again, using a probabilistic forecast of reservoir inflow and uncontrolled runoff. The differences between deterministic and stochastic optimal releases reflect hedging decisions that balance hydrologic risk. The difference in flood damage between deterministic and stochastic operation is one measure of the cost of hedging.

3.1 Operating Considerations

Flood control operation seeks to minimize flood damages by storing damaging inflow until it can be safely released. Beard (1963) characterizes good flood control operation as:

> "...releasing water whenever necessary at the highest practical rates so that minimum space need be reserved for flood control."

Throughout this analysis we assume that maximum discharge is the sole measure of flood damage. The duration of flows above flood stage is not considered. A rating curve and stage-damage function could be incorporated in the operating algorithms considered in this section in order to transform maximum discharge into expected damages. Such a transformation would guide operations by identifying stages at which incremental damages change. Using a stage-damage function, for example, an operator could recognize the level at which a flood plain is likely to be inundated. For areas with limited flood plain development but significant structural development above the flood plain, operations may allow inundation of the flood plain (see e.g. Coomes (1969)). In this way flood storage could be reserved for flows that would cause damage to the developed areas at higher elevations. Without loss of generality peak discharge is the only measure of damage considered, recognizing the value of incorporating the site specific non-decreasing stage damage function for all applications. Under this stage-damage assumption, the objective of flood control operation is to minimize the maximum stage. Lacking a particular stage-discharge function the goal of flood control operation is restated as minimizing the maximum discharge.

The damage caused by a current release will depend not only on the magnitude of that release, but also on the uncontrolled runoff downstream. Releases will be modified by channel routing and contribute to the downstream stage along with uncontrolled runoff. Concentrating on the maximum stage at a downstream damage center rather than the maximum release from the reservoir requires operating decisions that consider not only the effect of reservoir inflow, but also channel routing effects and the magnitude and relative timing of uncontrolled runoff downstream. For this reason, the timing of uncontrolled runoff critically determines the damage potential of a reservoir release. When uncontrolled runoff increases, reservoir inflow must be stored to prevent additional contributions to flood stage. When downstream stages are low, additional inflow can be released reserving reservoir storage for inflow with a higher potential for causing damage.

In this analysis operation on the rising limb of the flood hydrograph is distinguished from operation on the falling limb. On the rising limb, release decisions to store potentially damaging inflow must be balanced against the risk of unnecessarily committing storage that would be needed later in the flood. Balancing the benefits and risks of committing storage on the rising limb of the flood hydrograph is the critical operating tradeoff leading to successful flood control operation.

On the falling limb of the flood hydrograph, operation will still store inflow to prevent downstream stages from rising. In general the maximum volume of inflow held in storage will be found after the peak in inflow is observed. Falling limb operations are often more difficult than operating decisions on the rising limb. On the falling limb, rapid evacuation of the flood pool to reduce the risk associated with unexpected runoff must be balanced against both the magnitude and duration of flood stages at downstream damage sites. A non-trivial optimization problem can be formulated to guide efficient release of stored flood waters.

In this work the distinction drawn between rising limb and falling limb operation centers on the need to hedge against hydrologic uncertainty. On the rising limb of the flood hydrograph the dominant source of uncertainty is precipitation. A point is reached at which runoff-producing precipitation has ended and runoff can be accurately predicted through the use of an efficient network of precipitation and streamflow gages combined with meteorologic radar and hydrologic modeling. The technology to accurately gage watersheds in order to improve the accuracy of precipitation estimates and runoff forecasts is both available and cost-effective. Forecast accuracy can be further enhanced through the optimal design of precipitation networks (Hogan 1987). In this work appropriate forecast tools are assumed to exist to effectively eliminate hydrologic uncertainty on the falling limb of the hydrograph (after runoff producing precipitation has ended). All operating simulations evacuate the flood pool at the highest rate that will not increase the downstream stage.

Actual operations would likely seek to extend the drawdown period in order to lower the downstream stage from peak levels. This operating preference is easily accommodated in the current operating rules but is not central to the results presented here. For example, once runoff producing precipitation has ended, determination of drawdown releases can be treated as a separate operating problem. The constraining boundary conditions for drawdown are the volume held in storage and the maximum downstream stage when drawdown begins. Both are determined by hedging releases on the rising limb.

3.2 Deterministic flood control operation

If both reservoir inflow and uncontrolled runoff were known with certainty, a deterministic optimization problem could be solved to identify the optimal releases that minimize the maximum discharge at a downstream damage site. Such a problem has been formulated and solved for the Bloomington SPF shown in Figure 1. For the SPF, the maximum reservoir inflow of 91,400 cfs occurs at hour 53 and the maximum uncontrolled runoff of 51,000 cfs occurs at Luke, Md. at hour 51. With no flood control protection, the combined runoff hydrographs produce a maximum flood peak at Luke, Md. of 131,000 cfs at hour 53. Operation of the 36,000 AF flood control pool in Jennings Randolph Reservoir to reduce the flood peak at Luke is simulated.

If there were no uncontrolled runoff, the optimal operation of the flood pool would store all inflow above a prespecified threshold. Schwartz (1986) characterizes this regulation of the flood hydrograph as pure attenuation. A parametric operating rule described in Smith et al. (1987) achieves pure attenuation of the flood hydrograph by setting the release as

$$R_t = \mathrm{Min}\,\{Q_t\ ,\ R_{max}\}$$

where:

R_t = release rate for hour t
Q_t = reservoir inflow for hour t
R_{max} = threshold above which inflow is stored

Lowering R_{max} lowers the maximum discharge immediately downstream from the reservoir, increasing the volume of flood control storage used. Optimal deterministic operation sets R_{max} at that discharge that just causes the flood control pool to fill. In this way the full flood pool is used to minimize the maximum release.

As damage centers downstream from the reservoir are considered, flood control releases must account for the effects of uncontrolled runoff and channel routing in order to minimize the maximum downstream discharge. Using both the inflow and uncontrolled runoff hydrographs as well as a deterministic channel routing model, optimal deterministic reservoir releases will again fully utilize available flood control storage in order to minimize the maximum discharge at a downstream damage center.

The deterministic operating problem is solved for the Bloomington SPF using the existing allocation of 36,000 AF of flood control storage. The optimal releases and resulting Luke discharge are shown in Figure 2. The optimal use of flood control storage minimizes the maximum discharge at the downstream damage site (in the example considered here, Luke Maryland). Optimal releases account for the timing in both reservoir inflow and uncontrolled downstream runoff as well as delay and attenuation of releases due to channel routing. The optimal releases, shown in Figure 2 store the fraction of inflow that would contribute to a higher maximum discharge at Luke, while fully using 100% of the available flood storage during the routing of the SPF.

Several features of this deterministic optimal regulation are noteworthy. First, this optimal regulation strives to minimize only the maximum discharge at the downstream damage site. The duration of this maximum discharge is not considered. Ignoring the duration of the peak stage is not restrictive in considering these operating rules. The duration of the maximum flood stage could be shortened if a) the maximum stage was increased or b) the drawdown period on the falling limb of the flood hydrograph was extended. A reservoir operator is unlikely to choose to inundate developed areas, causing additional flood damage, in order to shorten the period of inundation. However, for damage sites with significant undeveloped floodplains, a decision could be made to increase flood stages, inundating the undeveloped flood plain, in order to reserve finite reservoir storage for later damaging flows. Coomes (1969) describes a three phase operating strategy for flood control reservoirs on the Kansas river. Using the stage damage curve, operating strategies shift from providing full protection against floods of moderate size, to planned inundation of cultivated bottom land, in order to provide maximum protection for the cities of Topeka and Kansas City during events with recurrence intervals in excess of fifty years. This type of strategy would be reflected in the use of a stage damage function to guide operations. The stage damage function is usually developed as a non-decreasing function of maximum stage. For this reason operations are, without loss of generality, evaluated based on the maximum discharge alone.

The second feature of the deterministic optimal regulation shown in Figure 2 is that the flood gates are never completely closed. Even during the peak of reservoir inflow and uncontrolled runoff, some releases are being made from flood control storage. These releases will inevitably contribute to the maximum discharge at Luke. If these releases were not made the Luke stage would be reduced during the peak of the runoff hydrographs. The extra storage used to withhold releases would not, however, be available on the falling limb of the flood hydrographs. The maximum discharge at Luke would actually increase on the falling limb as storage filled and inflow had to be released. Alternate release rules can store more inflow prior to the flood peak, reducing flood stages early in the flood event. The consequence of this short-term reduction in flood stage would be a premature commitment of flood storage, resulting in higher downstream discharges on the falling limb of the flood hydrograph.

The third feature of the optimal deterministic regulation is the time at which the flood pool fills. Even though the runoff peak occurs at hour 51, for the uncontrolled runoff, and at hour 53 for the reservoir inflow, the reservoir does not fill until hour 59. The optimal regulation to minimize the maximum damaging discharge reduces flood stage on the rising limb to the level that can just be sustained by filling reservoir storage during the storm. Even after the peaks of the runoff hydrographs have been observed and streamflow is clearly falling, inflow must be stored to prevent downstream discharge from rising. As Figure 2 shows, releases are less than inflow until hour 59 when reservoir storage fills. The deterministic regulation of the SPF reduces the Luke flood peak from 131,000 cfs to 76,000 cfs. This level of flood protection represents the largest reduction of the maximum discharge at Luke that could be realized through operation of the 36,000 AF of flood control storage in Randolph Reservoir.

In real-time operation, changing meteorologic conditions introduce uncertainty into the runoff forecast. The flood peak reduction achievable through the optimal deterministic operation of the flood control reservoir therefore represents an upper bound on the achievable flood benefits. The goal of real-time operation may be posed as achieving as great a proportion of these maximal benefits as possible. The following section considers stochastic real-time flood control operation. The need to hedge against hydrologic uncertainty is discussed in stochastic sequential decision problems. Real-time operation for the Bloomington SPF is simulated with a stochastic real-time operating rule. The hedging cost of hydrologic uncertainty as well as the characteristics of hedging operation are quantified by comparing stochastic and deterministic operating simulations.

3.3 Real-time Flood Control Operation

In this section the use of forecast information in stochastic reservoir operation is developed for the problem of real-time operation of a flood control reservoir. The probabilistic nature of streamflow and precipitation combined with finite storage capacity requires operating decisions that trade-off conflicting risks. Given the uncertainty of reservoir inflow during a storm event, reservoir operation must hedge flood control release decisions against estimation errors in both the timing and the magnitude of flood runoff. Each operating decision is a commitment of the limited storage that is available. This commitment of storage must balance the damages prevented by storing inflow against the risks of using limited storage inefficiently.

In real-time operation, the reservoir operator must make sequential release decisions, committing limited reservoir storage, based on imperfect probabilistic forecasts of runoff. There are two ways in which the operator can err in the face of uncertain inflow. If the magnitude of runoff is underestimated, the operator will commit too much storage too soon in the flood event. When the true flood peak is realized insufficient storage will be available. In this case the operator will be forced to release inflow, causing damages that could have been prevented. If the magnitude of runoff is overestimated, the operator will release the rising inflow, reserving storage for a flood peak which will not appear. In this case the operator will fail to store the true peak flow, again causing damages that could have been prevented.

The operator needs to make prudent operating decisions while hedging against hydrologic uncertainty. In this allocation decision, hedging implies balancing the short-term loss of benefits against much greater possible losses in the future. The unique characteristic of hedging is the certainty with which immediate, short-term costs will be incurred, in order to provide a level of insurance against uncertain future outcomes. Hedging decisions increase operating costs with certainty, in order to reduce the costs of risk taking. The real-time operating problem is a problem in optimal hedging.

The sequential hedging problem can be formulated as a stochastic optimization problem. At each hour the release is determined in order to optimize the allocation of a fixed volume of reservoir storage over the entire operating horizon, based on forecasts of runoff. Let the state of the reservoir system at any time t, be uniquely defined by the available volume of flood control storage, S_t, the most recent release, R_{t-1}, and the estimated parameter P_t of the current probability density of runoff, $Q(P_t)$. Then the expected cost of operation for the flood risk represented by the parameter vector P_t, starting in state $\{S_t, R_{t-1}, P_t\}$ is taken as

$$E[\ f(S_t, R_{t-1}, P_t)\ ] = \int_{P_t} f(S_t, R_{t-1}, P_t)\ dF_{P_t}(p_t)$$

in which E[.] is the expected value operator, $f(S_t, R_{t-1}, P_t)$

denotes the optimal cost (i.e. minimum flood damages) of flood control operations through the end of the operating horizon starting in state $\{S_t, R_{t-1}, P_t\}$, and the integral is taken over the estimated density of the parameter P_t.

If the immediate cost of making a release can be written as $c(R_t)$ (measured in units of storage) then the optimal release in time t will minimize the sum of the immediate cost plus the expected "cost-to-go" following the resulting probabilistic state transition. Then

$$E[f(S_t, R_{t-1}, P_t)] = \min_{R_t}\{c(R_t) + E[f(S_t - c(R_t), R_t, P_{t+1})]\} \tag{3.1}$$

in which we assume 1) the expected inflow under P_t is embodied in $c(R_t)$, completing the transition for the hydrologic component of the state vector and 2) the hydrologic forecast at each time step is uniquely parameterized by P_t. The optimal stochastic release is prescribed by solving problem (3.1) for R_t.

Each hourly operating decision incurs a cost $c(R_t)$ with certainty, in order to minimize the expected flood damages over the entire operating horizon. Incurring this hedging cost with certainty guarantees that the flood protection provided in real-time operation will be less than the level of protection possible from deterministic optimal releases. In return for bearing this certain cost, stochastic optimal operation will reduce the cost of risk taking by prescribing releases that are robust against the distribution of hydrologic realizations described by the density function for runoff, $Q(P_t)$), parameterized by P_t.

3.3.1 Operating with a Forecast

Instead of operating with a known description of the runoff hydrographs for the watershed, real-time operation must prescribe hourly releases based on the current runoff forecast. Forecasts are often described in terms of a single hydrograph, representing an expected value hydrograph. In this report the forecast is assumed to be described by a probability distribution of runoff hydrographs. To simulate stochastic real-time operation for the Bloomington SPF, the SPF hydrographs are described by a probability distribution. For the SPF flood hydrograph, both the magnitude and the timing of runoff are treated as random variables. The peak of the uncontrolled runoff at Luke is forecast only to within a 12 hour range and is equally likely to occur before or after the peak in reservoir inflow. Similarly the magnitude of both reservoir inflow and uncontrolled runoff is forecast with an assumed accuracy of $\pm 20\%$.

With this probabilistic forecast replacing the deterministic description of the runoff hydrographs, the operating strategy will change. Releases will still attempt to minimize the maximum downstream discharge, but these releases must now balance the reductions in downstream discharge that could be realized for some hydrologic realizations against the increased damages that would be caused for other, equally likely realizations. Stochastic optimal releases lead to incremental increases in damage producing stages in the short term in order to limit the magnitude of severe damages from flow levels that could arise later in the event.

3.3.2 Constant Hedging

To demonstrate the way in which hedging against hydrologic uncertainty changes operating decisions, the deterministic and stochastic optimal release decisions for the Bloomington SPF are shown in Figure 3. The stochastic optimal decisions are based on the joint distribution of timing and magnitude of peak runoff described above. Each hourly decision is determined by solving a stochastic optimization problem defined by specifying current reservoir storage and recent releases currently in channel storage, as initial conditions. For the operation shown in Figure 3 the forecast is assumed to be constant throughout the event. In other words the probabilistic component of the state vector in problem (3.1), P_t, is assumed to be constant for all time periods. The constant level of forecast uncertainty is analogous to the use of a constant safety factor in a target seeking operating rule. Hedging against a constant level of uncertainty is equivalent to constant hedging.

3.4 Stochastic vs Deterministic Operation

The differences between the optimal deterministic and optimal stochastic releases reflect the relative dominance of three types of hedging throughout the flood event. From hour 41 through hour 47 the optimal stochastic operation prescribes lower releases than the deterministic optimal releases in order to hedge against the possibility of an early peak in uncontrolled runoff. The cumulative difference in release represents the storage required to hedge against the possibility of these hydrologic realizations. In this case the optimal stochastic releases provide protection against early peaks in uncontrolled downstream runoff by committing flood storage. Early in the flood event, flood control storage is relatively abundant and the expected benefits from providing "early peak" protection are greater than the expected cost from a reduction of flood control storage.

The hedging decision is to commit storage with certainty rather than risk contributing to flood damages. As flood control storage is committed the expected damages from later, large magnitude flood peaks increase as well. From hour 47 through hour 57 the optimal stochastic releases are actually greater than the deterministic optimal releases. This occurs for two reasons. First, the deterministic releases are based on perfect knowledge of the runoff hydrographs. There is no doubt in deterministic operations about either the magnitude or the time of the inflow peak. The optimal deterministic releases decrease as uncontrolled runoff increases in order to prevent an increase in the maximum downstream discharge. In contrast, the stochastic optimal releases are based on a probabilistic forecast that includes the true flood hydrograph, as well as runoff events of considerably larger magnitude. Although flows are rising, less inflow is stored in order to provide protection against the larger flood peaks that the probabilistic forecast indicates are possible.

The optimal stochastic releases are also greater than the optimal deterministic releases due to the initial storage conditions. To successfully hedge against larger flood peaks, more inflow must be released in order to reserve flood control storage. This need is exaggerated by the higher initial storage levels created by the hedging decisions made for early flood peaks. Hedging against high flows seeks to minimize the consequences of both forecast risk as well as risk due to the reduced volume of available flood storage. The hedging decision is to release damaging inflow rather than risk having inadequate storage if a larger, later flood peak were realized.

Hedging against a third type of risk is demonstrated in the releases made on the falling limb of the hydrograph from hour 58 through hour 70. As discussed earlier, the maximum use of flood storage generally occurs on the falling limb of the hydrograph as inflow is stored to prevent secondary rises downstream. In order to prevent the downstream discharge from rising, while evacuating storage as rapidly as possible, the optimal deterministic releases gradually increase as uncontrolled runoff falls. The optimal stochastic releases hedge against higher discharges on the falling limb by reducing releases in order to assure that downstream discharges do not rise. The hedging decision is to store water in order to reduce the probability of a rise in stage downstream. The cost of these hedging releases is to increase the time required to evacuate the flood control pool.

The differences between deterministic and stochastic optimal releases reflect the way in which the stochastic operating rules respond to changes in relative operating risks. The need to hedge against hydrologic uncertainty indicates the fundamental difference between the meaning of an optimal decision in a stochastic and a deterministic environment.

In stochastic sequential decision problems such as the real-time problem of flood control operation, the optimal stochastic decision balances the risks associated with hydrologic uncertainty and the current state of the reservoir system. The cost of this hedging is the inevitable increase in damaging flows downstream. Post hoc, the optimal stochastic decision does not in general minimize flood damage. It does however, provide protection against a much wider range of flood events during the storm than the deterministic optimal releases. The intangible cost of risk-taking is reduced. The price of this reduction is a firm increase in incremental flood stages. The increase in flood stage represents the cost of uncertainty. The cost of hedging is the price paid to provide protection over the range of conditions described by the hydrologic forecast. If the forecast can be improved to increase accuracy and reduce uncertainty, the cost of hedging should be similarly reduced. The next section considers the role of forecasts and hydrologic uncertainty in real-time flood control operation.

4. Hydrologic Uncertainty in Flood Control Operation

Hedging in real-time flood control operation must be a part of real-time release decisions due to the uncertain nature of runoff during a flood. Qualitatively, the comparison between stochastic and deterministic operation in section 3 identified hedging releases dominated by three different sources of uncertainty, each of which resulted in a different type of hedging release. In this section the three sources of uncertainty identified in section 3 are considered in the broader context of real-time forecasting. The relative dominance of each source of uncertainty changes throughout the operating horizon, and is reflected in reservoir performance. The change in each type of uncertainty and the resulting need for hedging can be characterized in terms of the change in the probability distribution describing runoff. The importance of the rate of change of uncertainty in flood control operation is considered and two ways of quantifying this change are developed. The representation of each of these changes in the probabilistic forecast used to drive the real-time operating rule is considered and representative tools for simulating improvements in forecast skill are developed.

4.1 Hydrologic uncertainty

In this section, three changes in the forecast of runoff are considered: 1) changes in uncertainty on the falling limb of the hydrograph, 2) changes in the probability of the mode of the runoff distribution and 3) changes in the lower bound of the conditional runoff distribution based on current observations.

On the falling limb of the flood hydrograph, runoff is not known with certainty. The real-time predictions of the transformation of precipitation to streamflow will include errors in the point measurement of precipitation, errors associated with the extrapolation of point measurements to areal estimates of precipitation, error in the modeling of both infiltration and runoff, as well as model error and calibration error. The predicted falling limb hydrograph is viewed as a function of random variables and is described by its probability density function. While this general description of forecast error will be true for the entire flood hydrograph, it is on the falling limb of the hydrograph that uncertainty is dominated by uncertainty in the transformation of precipitation to runoff. In contrast, uncertainty on the rising limb is dominated by uncertainty in the spatial and temporal distribution of precipitation that has yet to affect the basin.

The uncertainty in the transformation of precipitation to runoff can be reduced by intensive instrumentation and watershed modeling (see e.g. Hogan(1987), Eagleson(1967), Grayman and Eagleson (1971) Hendrick and Comer (1970), Zawadzki (1973)). A sufficiently dense precipitation network, combined with meteorologic radar and an intensive watershed modeling effort could reduce the runoff uncertainty on the falling limb to a level that would be negligible for purposes of reservoir operation.

In flood control operation, uncertainty on the falling limb gives rise to hedging decisions regarding when to commit the remaining volume of flood control storage, and how rapidly flood storage can be evacuated. Reducing the uncertainty on the falling limb will speed drawdown by eliminating releases that hedge against the possibility of higher flows on the falling limb. Peak discharge, the measure of flood damage used in this report, will not in general be affected by a reduction in uncertainty on the falling limb. The maximum downstream discharge will be determined by the hedging releases on the rising limb of the hydrograph. Reduced falling limb uncertainty will ensure that the maximum discharge is not exceeded and will allow more efficient drawdown of the flood pool. In the operating simulations performed in this work, it is assumed that sufficient effort has been expended on gaging and modeling the entire watershed so that a time is reached on the falling limb at which runoff can be predicted with negligible error.

4.2 Improved forecast skill

As the set of observations used to prepare successive forecasts of runoff grows, the expanding information set should lead to a more accurate forecast, converging to the actual runoff, later in the flood event. This type of forecast skill can be characterized by the rate of reduction in the variance of the forecast as successive observations are made. Consider the real-time estimation of the total depth of precipitation over a watershed during a single storm. The total precipitation depth is a random variable. As an example, assume the forecast of precipitation is normally distributed with a mean equal to the true value of 100 mm and a variance of 225 mm^2. For this illustration, assume successive observations of the active, moving storm system are made, and point measurements are continuously collected from a network of precipitation gages. After the storm has passed we assume the realization of storm depth can be known with negligible error. The effect of sequential observations is to reduce the variance of the estimate of storm depth from 225 mm^2 to nearly 0 over the observation period. Without loss of generality we assume forecast skill will lead to a reduction in variance at an exponential rate. Then the estimated variance at any time t is described as :

$$s(t) = \exp\{ \ln[s(t0) \{1-(t-t0)/(T-t0)\}] \}$$

where:

$s(t)$ = the estimated variance at time t
$s(t0)$ = 225 mm^2
$s(T)$ = 0

As the successive density functions shown in Figure 4 suggest, the estimate of the total storm depth will become more accurate as the set of observations grows. At time T the storm has been completely observed and the variance of the estimated precipitation depth goes to zero. The only source of variance in the estimated depth after the storm has ended is the measurement error of the precipitation gage network, which we assume is negligible for our purposes. The reduction of forecast variance is one form of forecast skill considered in the operating simulations.

4.3 Bayesian update

Considering a forecast as a probability density function, the previous section suggests one way forecast skill improves real-time forecasts is by reducing the variance of the forecast. In this section, the view of a forecast as a density function is extended to consider the current forecast as a conditional density function, based on prior knowledge and experience with the system, and the current set of observations. An example using Bayes theorem is developed to suggest the way in which successive observations can be used to update the current conditional density, forming the forecast. In the context of the precipitation estimation example of section 4.2, a second example demonstrates that a Bayesian procedure will not only reduce the variance of the forecast density, but also truncate the density. The reduction of variance, and truncation of the conditional density, are emphasized as general characteristics of sequential forecasting. Alternate forecast algorithms will achieve these changes in the forecast density to varying degrees. The common characteristics of sequential forecasts, variance reduction and truncation of the conditional forecast density, are the features of real-time forecasting of greatest value in real-time reservoir operation.

4.3.1 Bayesian Theorem a Numerical Example

Bras and Rodriguez-Iturbe (1985) discuss hydrologic forecasting by considering the function

$$F(x(t_n)\,|\,x(t_{n-1}),x(t_{n-2}),\ldots)$$

defined as the conditional cumulative density of the random variable of interest at time t_n, $x(t_n)$, conditioned on the past values of the random variable. They note that a hydrologic forecast algorithm would attempt to calculate the conditional density

$$f(x(t_n)\,|\,x(t_{n-1}),x(t_{n-2}),\ldots) = \frac{f(x(t_n),x(t_{n-1}),x(t_{n-2})\ldots.)}{f(x(t_{n-1}),x(t_{n-2})\ldots.)}$$

in which the denominator contains all of the information in past observations. This joint density describes the likely state of the system at time t_{n-1}. The forecast is developed as the density of $x(t_n)$, conditioned on the density of the state of the system.

Bras and Rodriguez-Iturbe (1985) go on to note that in practice, neither the state of the system nor the observation of the system output is perfectly known when a forecast must be made. For this reason the process of sequential forecasting will consist of alternate steps of filtering and prediction. The filtering step uses all prior observations to estimate the conditional density of the system state at time t_{n-1}. The prediction represents the conditional density of the system output, given the prior density of the system state.

A simple example is developed to illustrate the use of Bayes theorem in sequential estimation of a forecast distribution. Consider a reservoir operator concerned about runoff from an ungaged upstream sub-basin in the drainage area controlled by a flood control reservoir. Although there are no precipitation or streamflow gauges on the watershed, the operator is experienced and familiar with regional runoff characteristics and has estimated the average response time of the watershed. Following the work of Huff (1967), Grace and Eagleson (1967), Pilgrim et al. (1969) and Eagleson (1970), assume that analysis of the dimensionless mass duration curve for precipitation has allowed the operator to recognize that storms can be classified by determining whether the peak runoff occurs early in the storm or late in the storm, compared to regional average runoff.

The reservoir operator needs to know if the current storm peak is early, average, or late, compared to runoff over the remainder of the watershed, for which accurate telemetered hydrologic data is available. The uncertainty in the time of peak from the ungaged watershed will cause the operator to hedge reservoir releases in order to avoid making a release that compounds peak runoff. The operator's release decision will be guided by the probability distribution of time to peak for the runoff from the ungaged sub-basin, and it is this distribution that we attempt to estimate using Bayes theorem.

At each hour, for the ungaged basin, the operator wants to calculate

$$P\{t_{pi} | Q_t, Q_{t-1}, \ldots\ldots\} \qquad (4.1)$$

the conditional probability of the i-th time of peak given the observed flows to date, in which $i = \{1,2,3\}$ corresponds to early, average, or late peak runoff respectively. Although the sub-basin is ungaged, assume the operator has a friend who lives on the sub-basin, and each hour the operator makes a phone call, and his friend tells him the stage of the river is either low, medium, or high. The operator's problem is to use these "observations" of the river to update the conditional distribution of the time of peak for the ungaged sub-basin. (These qualitative observations are similar to the work of Tsonis (1987) in which GOES satellite imagery cannot be used to estimate rainfall rates, but the identification of "light-moderate" and "moderate-heavy" precipitation areas is possible).

From his past experience, the operator begins his operation with no knowledge of the state of the flood. The initial probability of the time of peak (4.1) is therefore assumed to be uniformly distributed. Early in the storm, the observer will describe the river as low, regardless of the time to peak. As the river rises, the time at which the imperfect description of the stage becomes medium or high allows the operator to modify his judgement as to the time of peak runoff for the sub-basin. The operator would expect an early peak to have a high stage at the current time, an average storm would have a medium stage at the current time, and only a late storm would produce a low stage. If the system could be perfectly observed, one error-free observation would be sufficient to determine the time to peak of the storm. This uncertainty of the observation is described by the likelihood $P(t_p|Q)$.

For this example assume the reservoir operator's experience leads him to the prior likelihood, $P(Q|t_p)$ described by the following matrix:

Relative Time of Peak			
Early	Average	Late	Observed Stage
0.24	0.32	0.53	low
0.30	0.41	0.27	medium
0.45	0.27	0.20	high

In other words, if the runoff peak was actually later than average, the operator believes there is a 53% chance that the observer will report the stage as low. Due to the uncertainty of storm characteristics as well as the subjectivity of the observer, the operator believes that if the storm has an average time of peak, the observer could still report the stage as low with a probability of .32, and with probability .24 the storm could have an early peak and still result in the observer characterizing the stage as low in the current hour.

To account for measurement and observation error the operator uses a sequential Bayesian procedure in each hour to update and refine the conditional probability distribution used to determine reservoir releases.

At every time t the operator wants to calculate

$$P(t_{pi}|Q_t) = \frac{P(t_{pi},Q_t)}{P(Q_t)} = \frac{P(Q_t|t_{pi})P(t_{pi})}{P(Q_t)} = \frac{P(Q_t|t_{pi})}{\sum_{i=1}^{3} P(Q_t|t_{pi})P(t_{pi})}$$

At each hour the operator has a prior distribution for the time of peak, $P(t_{pi})$, and the current likelihood matrix, as above, containing $P(Q_t|t_{pi})$. When the observer reports the current stage, the operator reevaluates the probability of the time of peak (to make the current reservoir release) by using $p(t_{pi}|Q)$ as the posterior estimate of $P(t_{pi})$. The operator may also have a procedure to revise the likelihood matrix $P(Q_t|t_{pi})$ as successive observations in the storm make one of the alternatives more likely. This procedure is demonstrated in several examples.

EXAMPLE 1. Assume the storm has an average time of peak, for which "medium" stage height observations would be observed. Let the observer correctly observe the stage height as medium in each hour, and use the likelihood matrix above without modification.

Starting with a uniformly distributed prior distribution $P(t_{pi})$ the calculations for one medium observation of stage are presented below.

Prior Probability of
Relative Time of Peak

Early	Average	Late	
0.33	0.33	0.33	
Likelihood of Observation			$P(Q_t\|t_{pi})$
0.24	0.32	0.53	low
0.30	0.41	0.27	medium
0.45	0.27	0.20	high
Bayesian Posterior		$P(t_{pi}\|Q_t)$	
0.22	0.29	0.49	low
0.31	0.42	0.27	medium
0.49	0.29	0.22	high

The effect of the first observation can be seen by comparing the prior distribution of t_p to the posterior that would result if the observer reported each of the three possible stage heights. For this example, assuming the observer correctly reports a "medium" stage height, the prior distribution of t_p for the following hour will be

$$P(t_{pi})=\{.31,.42,.27\}.$$

The prior probability $P(t_{pi})$ is recursively calculated using the constant likelihood matrix above, assuming the observer correctly estimates the stream stage to be medium at each hour. As Figure 5 shows the probability of the (assumed) correct estimate of the time of peak grows to 0.8 in 6 hours. This increase in the probability of the "true" realization of the forecast distribution is a crucial attribute of recursive forecasting that will be utilized in simulated flood control operation in sections 5 and 6.

EXAMPLE 2. In this example, again assume the storm has an average time of peak but now the observer incorrectly judges the stage to be high for the first two hours before correctly identifying the stage as medium for all subsequent hours. The likelihood matrix is again assumed fixed.

The operator's prior distribution of t_p in the first hour after the "high" observation is taken as

$$P(t_{pi}) = (.49,.25,.26).$$

In the second hour this increases to

$$P(t_{pi}) = (.66,.17,.17)$$

The incorrect identification of the stage in the first two hours has skewed the distribution of t_p. Subsequent observations result in the shift in probability mass to expect an average peaking hydrograph. The effect of the initial bad guess is to reduce the rate at which the probability of the (assumed) true realization of t_p grows. Figure 6 shows the change in probabilities, $P(t_{pi})$ for the first 7 hours of the forecast. Although the probability of an average time to peak is only 0.17 at hour 2, by hour 7 this probability has increased to over 0.7.

EXAMPLE 3. In this example we again assume the storm has an average time of peak, and the observer correctly identifies the stage as medium for each of the first 6 hours. As the observer continues to report medium stages, the likelihood $P(Q|t)$ is updated. After the third consecutive report of a medium stage, the probabilities of a "correct" stage observation given the "true" peak time increase 1-2% per hour. This could reflect an increased confidence in the observer's accuracy. By reducing the observational uncertainty the effect of a single observation increases, approaching the case of perfect certainty, in which one observation of stage is sufficient to uniquely identify the time of peak for the storm.

The Likelihood matrix for each hour is presented below.

	Early	Average	Late	
	Likelihood of Observation- P{Q|t}			
hour 0	0.24	0.32	0.53	low
	0.30	0.41	0.27	medium
	0.45	0.27	0.20	high
	Likelihood of Observation- P{Q|t}			
hour 1	0.24	0.32	0.53	low
	0.30	0.41	0.27	medium
	0.45	0.27	0.20	high
	Likelihood of Observation- P{Q|t}			
hour 2	0.24	0.31	0.55	low
	0.29	0.42	0.26	medium
	0.47	0.27	0.19	high
	Likelihood of Observation- P{Q|t}			
hour 3	0.23	0.30	0.56	low
	0.29	0.43	0.25	medium
	0.49	0.26	0.19	high
	Likelihood of Observation- P{Q|t}			
hour 4	0.22	0.30	0.58	low
	0.28	0.45	0.24	medium
	0.50	0.26	0.18	high
	Likelihood of Observation- P{Q|t}			
hour 5	0.22	0.29	0.59	low
	0.27	0.46	0.24	medium
	0.51	0.25	0.18	high

Using the updated likelihood, the probability of an average time to peak grows faster than for example 1, reflecting the increased confidence in the observer's reports. Figure 7 shows the change in the probability distribution through hour 6 using the updated likelihood. Figure 8 compares the effect of the updated likelihood on the probability of both the early and the average peak estimate. Small changes in the likelihood produce a significant increase in the mode of the forecast distribution.

The problem posed in this section was formulated to demonstrate the way in which the information content of imperfect, sequential observations, could be incorporated in a conditional probabilistic forecast setting using Bayes theorem. In real-time forecasting for real-time reservoir operation, the information state, the observation network, and the quantity of real-time data would be significantly more complex. No attempt is made to develop an operational forecast system in this report. Rather the main characteristics that any sequential forecast system will have are qualitatively identified in order to develop synthetic forecast series to simulate real-time operation. The following section suggests a partial framework in which an operational forecast would be developed.

4.3.2 Sequential Estimation

In the section 4.2 , the reduction in forecast variance was attributed to an increasing set of hydrometeorologic observations that could include meteorologic radar and atmospheric measurements as well as precipitation and streamflow gage readings over the basin. In this section the estimation of total precipitation depth during a storm is considered by estimating the distribution of the sequence of hourly precipitation values. If the storm is represented as a sequence of hourly precipitation depths, then the random variable X(t) is the depth of precipitation over the basin in hour t and the storm can be represented by the sequence

$$X = \{X(1), X(2), \ldots\ldots\ldots\ldots X(T)\}.$$

The sequence X is the realization of a stochastic process parameterized by the parameter vector P, with density $f_P(p)$ and can be described through the joint distribution of X and P,

$$f_{XP}(x,p).$$

From the sequence of observations $\{x(1), x(2), \ldots.x(t)\}$, an estimate of a function of X, G(X) is derived. This function may be total storm depth, D

$$D = \sum_{i=0}^{T} X(i).$$

The observed history of the process at time t, H(t), is used to estimate E{G(X)} = E{D}, the expected value of total precipitation depth, in order to use the density of D as the operating forecast. The sequence of hourly observations will therefore generate a sequence of expectations,

$$\{E_1(d), E_2(d), \ldots.\ E_T(D)\}$$

where:

$E_j(D)$ is the expectation of total storm depth at time j, and $E_j(D)$ will be based on H(j), including $X_j = \{ X(1), x(2), \ldots.. X(j) \}$.

Each estimate of X is described by the marginal density of X,

$$f_X(x) = \int f_{XP}(x,p)\ dp.$$

Prior to observing the partial realization of X, X_t, where

$$X_t = \{\ x(1),\ x(2),\ .....\ x(t)\ \}$$

Assume the prior märginal density $f_X(x)$ had been calculated. From the new observation, x_t, the prior marginal density of X will be updated to yield the posterior marginal density of X,

$$f'_X(x) = \int f_{X|P}(x|p) f'_P(p)\ dP$$

where $f'_P(p)$ denotes the posterior density of P.

The observation x_t leads to an updated estimate of the marginal density of P, the parameter vector that drives X. From this posterior density for P the posterior density of X is calculated.

The posterior density for P, $f'_P(p)$ can be calculated using Bayes Theorem,

$$f'_P(p) = \frac{f_{X|P}(x|p)\ f_P(p)}{\int f_{X|P}(x|p)\ f_P(p)\ dP}$$

which is simply the product of the likelihood of x and the prior of P, normalized by the prior marginal density of X.

As an example, consider the estimation of total precipitation during a storm, from successive hydrometeorologic observations. Considering the unbiased estimate of the mean of the normally distributed storm total to be 100 mm, the density function will be updated using the observed storm total. The density progressively becomes truncated through the storm. Initially there is some probability of a precipitation total as low as, say 87 mm. This probability is just

$$\int_{-\infty}^{87} f_D(v)\,dv$$

Where $f_D(v)$ represents the density for the accumulated precipitation D. After a partial storm total has been observed to exceed 87 mm, the density for all values less than 87 mm must be zero. The conditional density function for storm total, given an observed total of at least 'm' mm is therefore

$$f_D(D) = 0 \qquad D \in [-\infty, m]$$

$$= K\, f_D(v) \qquad D \in (m, \infty]$$

Where K is a normalizing factor ensuring the integral of the density is unity. K is therefore seen to be

$$[1 - \int_m^{\infty} f_D(v)\, dv\,]^{-1}$$

The density for total precipitation, assumed to be normally distributed with a mean of 100 mm and variance of 100 mm^2 is shown in Figure 9, along with the truncated, conditional density resulting from observations of precipitation of 87 and 92 mm. In general, another manifestation of the reduction of forecast uncertainty resulting from the incorporation of successive observations, will be truncation of the forecast density.

In general, we would like to consider the joint density of the parameter P and the sequence X as

$$f_{X_0, X_1, \ldots X_T, P}(x_0, x_1, \ldots x_T, p)$$

A forecast of the next hourly rain total could be described through the calculation of the conditional density of X_{t+1} given X_t, $X_{t-1}, \ldots,$ X_0, P

$$f_{X_{t+1} | X_t, X_{t-1}, \ldots, X_0, P}(x_{t-1} | x_t, x_{t-1}, \ldots, x_0, P)$$

An iteration of estimation begins with the prior density of P, $f_P(p)$ and the likelihood of X_0,

$$f_{X_0 | P}(x_0 | p)$$

After observing x_0, the prior density of P and the likelihood of X_0 is used to calculate the posterior density of P, $f'_P(P)$,

$$f'_P(P) \propto f_{X_0 | P}(x_0 | P)\, G(p)$$

More generally, we are interested in the forecast described by the joint density of $X_0, X_1, \ldots, X_T, P$. Precipitation realizations through time t produce observations of $x_0, x_1, \ldots, x_t$. The remaining forecast requires the calculation of the conditional density of X_{t+1}, $X_{t+2}, \ldots X_T$, based on the observations to date:

$$f_{X_{t+1},X_{t+2},\ldots X_T \mid x_0,x_1,\ldots x_t,P}(X_{t+1},X_{t+2},\ldots X_T \mid x_t,x_{t-1},\ldots x_0,p)$$

This is proportional to the integral of the joint density of $X_t, X_{t+1}, \ldots X_T$ and P, taken over P, with the observed values of $x_0, x_1, \ldots x_t$ taken as parameters.

The joint density can be written as the product

$$f_P(p)\; f_{X_0|P}(x_0|P)\; f_{X_1|X_0,P}(x_1|X_0,P) \;\ldots\; f_{X_T|X_{T-1},X_{T-2},\ldots X_0,P}(x_T|X_{T-1},X_{T-2},\ldots X_0,P)$$

which is the product of the prior density of P and the joint density of $X_0, X_1, \ldots X_T$ conditioned on P. With successive observation of the X_i's, the joint density would be recalculated, giving the current forecast as the conditional joint density of future precipitation.

Summary

The effect of successive observations in the sequential estimation process is two-fold: updated estimates reduce the variance of the forecast and truncate the forecast density, bounding it from below. In this report the reduction in variance is represented by the rate at which the probability of the mode of the forecast distribution increases. The truncation of the forecast distribution leads to a redistribution of probability mass. This change is quantified by the change in negative entropy during the forecast period.

These characteristics of change in the forecast distribution are simulated in the forecast described in section 3.3. The following section incorporates simulated forecast skill with real-time operating rules to quantify the tradeoffs between storage reallocation and operating improvements.

5. SIMULATED FLOOD CONTROL OPERATION IN JENNINGS RANDOLPH RESERVOIR

In this section the effect of both storage reallocation and forecast skill improvements is analyzed through the simulation of real-time flood control operation of Jennings Randolph Reservoir. The stochastic operating rule is used to regulate the Bloomington Standard Project Flood using the existing 36,000 AF flood control pool in Randolph Reservoir. A probabilistic runoff forecast, with no change in forecast skill during the flood event (described in section 3.3) drives the operating rule. This constant level of forecast uncertainty is comparable to the level of forecast skill that would be available from regional forecast products such as the Quantitative Precipitation Forecasts that are prepared by the National Weather Service. Operating with constant forecast uncertainty is analogous to the use of a constant safety factor in a target seeking operating rule. A probabilistic forecast is identified that produces operating results comparable to the Corps of Engineers' design operation for the standard project flood.

Using this forecast, real-time operation is again simulated for reallocations of flood control storage of 5,600, 10,300, and 12,200 AF. The reallocation of storage with no change in operating rules leads to an increase in maximum downstream discharge. The operational effects of two types of forecast improvements are then simulated to quantify the value of improved forecast skill in supporting flood control storage reallocations. The value of optimal hedging using stochastic real-time operating rules is demonstrated. The allocation of reservoir storage, the choice of real-time operating rules, and real-time forecast skill jointly determine the operating costs and achievable benefits from a reallocation of reservoir storage. While improved forecasts should improve operations, the full information content in probabilistic real-time forecasts can only be realized through the use of stochastic real-time rules.

5.1 Reallocation with no forecast skill

Using a static probabilistic forecast, real-time operation of the existing 36,000 AF flood control pool in Randolph Reservoir was simulated for the Bloomington SPF shown in Figure 1. The probabilistic forecast is described by two random variables. The magnitude of the hydrographs for both reservoir inflow and uncontrolled downstream runoff is considered to be a random variable and is described by a probability distribution. Similarly the relative timing of the two hydrographs is not known with certainty. The difference, in hours, between the times of peak discharge for the two hydrographs is also treated as a random variable. The probabilistic forecast consists of the joint probability distribution of difference between time to peak, and magnitude of peak discharge. Initially the level of uncertainty associated with the forecasted runoff is high. This is represented by assuming that the initial distributions of both the magnitude and relative timing of peak runoff are independent and uniformly distributed.

Real-time operation for the Bloomington SPF using this highly uncertain probabilistic forecast reduces the peak discharge at Luke, Md. from 131,000 cfs to 96,334 cfs. The flood peak reduction achieved by the stochastic operating rule with this probabilistic forecast is nearly identical to the Luke peak of 97,500 cfs resulting from the SPF design operation (U.S. Army COE 1980). This constant probabilistic forecast, used with the stochastic operating rule is therefore used to describe the current operating practice for the flood pool in Randolph Reservoir. The effect of only reducing the size of the existing flood pool through a reallocation of storage is examined by operating the reduced storage volume using the constant probabilistic forecast to drive the stochastic operating rule.

Using the same probabilistic forecast, real-time operation is simulated for reallocated volumes of flood control storage of 30,400, 25,700 and 23,800 AF. The reduction in flood control storage results in an increase in maximum discharge at Luke as expected. The increased flood levels resulting from incremental reallocations of flood control storage are summarized in Table 1.

The effect of real-time operation of the reallocated volumes of flood control storage can be represented by the trade-off curve of storage and peak discharge shown in Figure 10. The level of flood protection resulting from reallocated storage is determined by the stochastic operating rules. These rules in turn, account for the change in storage capacity in determining the releases in each hour. This contrasts with the common target-seeking rules in which the hourly releases made under different storage allocations would be identical until the available storage was limiting.

These operating simulations quantify the loss in flood protection resulting from incremental reallocations of flood control storage in Randolph Reservoir, with no change in forecasting skill. The next section considers the extent to which real-time operation with improved forecast skill can compensate for the loss in flood protection resulting from the reallocation of flood storage.

5.2 Reallocation With Forecast Skill

This section considers the value of increased forecast skill during a flood event in reducing downstream flood peaks. Section

4 considered three qualitative improvements that could be realized in real-time flood forecasting: 1) reduced uncertainty on the falling limb of the flood hydrograph; 2) an increase in the probability of the mode of the forecast distribution as additional observations are made in real-time; 3) truncation of the forecast distribution as additional observations accumulate. The value of each of these types of reductions in hydrologic uncertainty is quantified through simulated operation on the Bloomington SPF.

5.2.1 Forecast Skill on the Falling Limb

Uncertainty on the falling limb of the flood hydrograph delays the drawdown of the flood pool. Hedging on the falling limb of the hydrograph stores inflow until storage can be evacuated without risking an increase in downstream stage. A dense network of precipitation gages combined with accurately calibrated hydrologic watershed models can eliminate most of the uncertainty on the falling limb of the runoff hydrograph. For the Bloomington SPF shown in Figure 1, the peak inflow to the reservoir occurred at hour 53. The hyetograph for the Bloomington SPF in Figure 11 shows the maximum hourly precipitation occurred at hour 42. By hour 48 effective precipitation had ended, and all measurable precipitation had ceased by hour 51. Real-time hydrometeorologic data collection could reduce the uncertainty associated with measuring the distribution and intensity of precipitation to negligible levels. The remaining uncertainty on the falling limb of the hydrograph would be largely associated with error from precipitation and streamflow gages and calibration and model error in the watershed model.

In contrast to uncertainty on the falling limb, uncertainty on the rising limb of the flood hydrograph is dominated by the structure, evolution, and motion of storm systems in the atmosphere. Put another way, uncertainty on the falling limb is dominated by the way in which precipitation is transformed into streamflow. Uncertainty on the rising limb is dominated by the uncertainty in forecasting precipitation that has not yet fallen. In all the stochastic operating simulations considered in this report, the runoff hydrograph is assumed to be known with certainty on the falling limb of the hydrograph. The hour at which falling limb uncertainty is assumed to be negligible is discussed in the following section.

5.2.2 Forecast Skill on the Rising Limb

Combining storm track images from meteorological radar with observations from precipitation and streamflow gages, would allow the estimate of runoff for the remainder of the flood event to be progressively refined. This sequential accumulation of information could be thought of as repeatedly estimating the parameters of the probabilistic forecast. The increasing set of observations effectively reduces the variance and truncates the forecast distribution as described in section 4. This section considers only the effect of reducing uncertainty by transforming the uniform forecast distribution into a unimodal distribution with reduced forecast variance.

The reduction in hydrologic uncertainty resulting from sequential hydrometeorologic observations is simulated by modifying the joint uniform distribution of the magnitude and relative time to peak of the probabilistic forecast. As the forecast improves during the storm, the probability associated with the mode of the joint distribution, corresponding to the SPF, increases. The value of increased forecast skill in real-time operation is relatively insensitive to static measures of forecast accuracy such as the rms error or the coefficient of prediction. In real-time operation, it is the rate of reduction of forecast error during the flood event that critically determines the level of flood protection that can be provided.

From the initial uniform distribution at time t0, the probability associated with the parameters of the actual runoff hydrograph will increase to unity at time T. The skill associated with improved parameter estimates due to successive observations is simulated by allowing the probability of the mode of the joint distribution of magnitude and time of peak runoff to increase at an exponential rate.

The probability of the mode of the forecast distribution at time t is calculated as:

$$P(t) = \exp \{\ln p(0) \; [\; 1 - (t-t0)/(T-t0)\,]\}$$

where:

- t0 is the hour at which forecast skill begins to improve compared to the uniform distribution.
- p0 is the initial joint uniform probability of all runoff realizations.
- T is the hour at which the uncertainty in the runoff hydrograph is assumed to be negligible.

Early in the storm the probability of the mode increases very slowly. At some time, T, the observations of precipitation and streamflow are sufficient to allow an accurate forecast of runoff; the probability associated with the true runoff hydrograph approaches unity. While the increase in the probability of the mode is small early in the storm, the forecast accuracy improves at an increasing rate near hour T.

Under this exponential model of improved forecast skill, the overall skill of the forecast process can be represented by the quantity $\{T-t_0\}$. As discussed above, the SPF is generated from the precipitation distribution in Figure 11. The SPF hydrograph peaks at hour 41 and is completely observable with insignificant error by hour 51. Under these circumstances it would be reasonable to expect to have an accurate forecast of runoff by hour 60. An increase in the rate of improvement of forecast skill is simulated by allowing the parameter T to take on smaller values. For the operating simulations considered in this section, three rates of forecast improvement are considered. The probability of the SPF realization begins to increase at hour 40 and increases exponentially to unity at hours 60, 53, and 50.

The accurate prediction of runoff for the Bloomington SPF at hour 60 is reasonable. By hour 60 (7 hours after the peak reservoir inflow and 19 hours after the peak precipitation), a sufficiently dense network of precipitation and streamflow gages could eliminate virtually all the uncertainty associated with the timing and magnitude of precipitation that had fallen. This would allow a calibrated hydrologic model to be verified for this storm and used to produce an accurate runoff forecast.

Runoff certainty at hour 53 results from improved forecasting skill. The required improvements are associated with real-time data collection, data processing, and hydrologic modeling. No skill in predicting the distribution and intensity of precipitation is required. Runoff forecast certainty at hour 50 indicates a level of accuracy in predicting runoff that allows the peak of both uncontrolled downstream runoff and reservoir inflow to be accurately forecasted before they are observed. While this represents a qualitatively higher level of forecast skill, the dominant source of uncertainty is still in the runoff and channel routing processes. The runoff-producing precipitation has essentially ended and accurate peak prediction will depend on the accuracy of the runoff and channel routing models used to transform the observed precipitation into streamflow. The increasing levels of forecast skill represented by the three values considered for the parameter T are judged to be reasonable, achievable levels of forecast skill with existing technology. Figure 12 shows the rate at which the probability of the mode of the forecast distribution increases for each level of forecast skill.

The distribution of the runoff parameters will determine the degree to which hedging occurs. Early in the storm when many hydrologic realizations appear equally likely, releases will be relatively cautious, hedging against the many flood hydrographs that appear to be nearly equally likely. As the probability of some of the more extreme events declines, release decisions will be made to mitigate the most likely flood hydrographs. In addition to changing the relative probabilities of alternate, feasible hydrologic realizations, the exponential change in probability described above will also change the information content of the forecast. To quantify this effect the negative entropy of the forecast, a measure of the information content of the probability distribution, is calculated as:

$$H = -\sum_{i=1}^{n} p(i) \log_2 p(i)$$

where p(i) is the probability of the i-th realization of the joint distribution of random runoff parameters. The change in negative entropy for each of the exponential forecast changes described above, is plotted in Figure 13.

Operation of the existing 36,000 AF allocation of flood control storage is again simulated for each of the three exponential rates of increase in forecast skill. The maximum discharge at Luke is compared for each forecast as well as the constant level of forecast uncertainty described in section 3. The operating results summarized in Table 2 show how real-time operating rules can incorporate forecast improvements during a flood event to improve reservoir operations.

The exponential reduction in forecast uncertainty at hour 60 significantly reduces the Luke flood peak by over 4000 cfs, or nearly 5% compared to operation with a constant level of forecast uncertainty. Although the relative improvement from using a forecast with increasing skill is significant, the marginal reductions in maximum discharge realized from a further marginal increase in the rate of improvement in forecast skill are small. The ability to accurately anticipate the flood peak is a major qualitative improvement in forecast accuracy compared to accurately identifying the peak runoff 7 hours after it has been observed. Accurate prediction of the peak would require an extended network of both precipitation and streamflow gages, as well as accurately calibrated hydraulic routing models for the major upstream tributary channels, for tributary runoff to be accurately estimated. The corresponding increased level of forecast skill, resulting in zero entropy at hour 50, provides an incremental reduction in Luke discharge of only 471 cfs, compared to the reduction in Luke discharge of 4,333 cfs achieved by using a forecast that does not achieve zero entropy until hour 60.

These results indicate the importance of hedging on the rising limb, in determining the downstream flood peak. The ability to reduce any hedging costs will provide additional flood control storage for damage reduction. The entropy of each of the forecasts is similar early in the event, when the hedging cost is high. For this reason the, incremental improvement in flood protection achieved from any improvement in forecast skill is much greater than any further marginal improvement resulting from increasing the rate at which forecast skill improves.

For all levels of flood control storage considered, the use of a dynamically improving forecast significantly reduces the maximum Luke discharge, compared to the constant probabilistic forecast. The incremental reductions in Luke discharge resulting from a further incremental increase in the rate of forecast improvement, are relatively small. Forecasts that can accurately predict the runoff hydrograph early in a storm will be most valuable for reservoir operation, and most difficult to refine. The operating simulations summarized in Table 2 indicate that the stochastic operating rule developed for flood control operation of Randolph reservoir provides the greatest marginal flood peak reduction from the forecast improvements that are most easily implemented.

The value of real-time forecast improvements becomes more clear when comparing the operating results for different storage allocations. Reducing flood control storage increases peak downstream discharge. Real-time operation using improving forecasts can reduce the peak downstream discharge. Using stochastic operating rules with dynamic real-time forecast improvements, the increase in flood damages resulting from a reallocation of flood control storage can be reduced and in some cases eliminated. The operating results summarized in Table 2, indicate that the increased flood damages resulting from a reallocation of flood control storage from 30,400 AF to 25,700 AF are more than offset through the incorporation of real-time flood forecasting in real-time operation. A reallocation of the current flood control pool of 36,000 AF to 30,400 AF with no change in the hydrologic forecast driving operations would result in an increase in the maximum discharge at Luke of over 5000 cfs. With the increase in forecast skill represented by the convergence to certainty at hour 50, the reallocated flood pool of 30,400 AF could be operated to produce an increase in the maximum Luke discharge of less than 150 cfs. The use of stochastic operating rules and real-time forecasts can increase the range of benefits available from existing projects, and significantly increase the range of feasible alternatives in multipurpose reservoir operation.

The value of a storage allocation in a multipurpose impoundment cannot be evaluated independently of the operating rules that will determine the release decisions. In real-time operation, the optimality of releases must be judged in terms of an optimal balance between hedging costs and hydrologic risks. Real-time forecasting is the key to successful real-time reservoir operation. Efficient operation, in turn, determines the value of an allocation. Reservoir design, or allocation, is inexorably linked to reservoir operation. Superior operation is in turn inseparably linked to skillful real-time forecasting. The value of any reallocation of reservoir storage must be judged with respect to the hydrologic forecasts used to drive stochastic, real-time, storage dependent operating rules.

In addition to the increased information content of the forecast resulting from improved parameter estimates, the information content of the forecast is also improved by estimates that truncate the conditional forecast distribution as the storm is observed. Successive observation allows the assignment of zero probability to some realizations of the runoff distribution as discussed in section 4.3. For example, if the inflow is still rising at hour 50, realizations that called for peak runoff at earlier hours are no longer possible. The elimination of the need to hedge against some hydrologic realizations is examined in the following section to suggest the value of a Bayesian estimate of the forecast distribution. The reduced probability of some runoff realizations redirects hedging decisions in order to protect against the most likely risks. A truncated forecast distribution describes the foregone need to hedge against these realizations.

5.2.3 Bayesian update

The value of reduced hedging costs due to truncation of the forecast distribution of runoff is quantified by simulating real-time operation for the Bloomington SPF for the four storage allocations considered above. The operating results are simulated using the initial probabilistic forecast with a constant level of uncertainty and a constant forecast. A single update at hour 43 eliminates the need to hedge against early peaks based on the observed uncontrolled runoff of only 9,220 cfs. The results shown in Table 3 indicate the substantial reductions in downstream flood peak achieved by reducing the need to hedge. The additional storage made available by eliminating the need to hedge against early peaking runoff events is used to achieve significant incremental reductions in downstream peak discharge for all storage allocations.

In this section the effect of early hydrometeorologic observations in reducing the need to hedge is examined and shown to offer significant operating improvements. The increased forecast information content considered earlier assumed the probability of the more extreme realizations of flood producing runoff became small through time. In this section observations are considered that would allow the need to hedge against some hydrologic realization to be eliminated. The consequences for such changes can be valuable when the hedging costs (in terms of storage) are significant.

The Bloomington SPF in Figure 1 is again considered, with both the magnitude and time of peak considered to be random variables following the joint uniform distribution used in the initial stochastic simulation in section 3.3. For this uniform distribution, some extreme realizations create the possibility of a peak in uncontrolled downstream runoff in excess of 80,000 cfs, as early as hour 45. These early peaks appear possible in the initial flood forecasts and dominate early releases as discussed in section 2. As sequential release decisions are made and both the inflow and uncontrolled downstream runoff in Figure 1 are observed, some of the large magnitude/early peak realizations can be ruled out. For example, the initial distribution admits an uncontrolled runoff hydrograph with runoff of 39,492 cfs at hour 40, peaking at hour 45. For this realization, the uncontrolled runoff expected at hour 43 would be 61,638 cfs. As the true uncontrolled runoff shown in Figure 1 is observed and inflow at hour 43 is observed to be only 11,526 cfs, some of these extreme realizations can be ruled out, making the remaining hydrographs more likely.

Routing of observed precipitation in real-time provides a lower bound on both the magnitude and time to peak of the flood hydrograph. For example, at hour 43 the hyetograph in Figure 11 indicates the maximum hourly precipitation depth has just been realized at hour 42. With this observation, the peak runoff will not occur for at least 5 to 6 hours (the concentration times for the basins). The probability of runoff hydrographs with earlier peaks must therefore be zero, effectively truncating the distribution of runoff. In addition to an increase in the probability of the mode of the forecast distribution discussed above, the distribution of runoff is also truncated and bounded from below as the conditional distribution is updated with real-time observations. As the probability of early peaks becomes low, hedging against these events that dominated operations in hours 40 - 46, is reduced and finally eliminated. Obviating the need for this hedging makes reservoir storage available for larger, later peaks, increasing the level of flood protection that can be provided.

To simulate this effect in a simple way the probability distribution used to drive operations is truncated at hour 43. The hyetograph in Figure 11 indicates that maximum precipitation occurred at hour 42. The three inch precipitation total in hour 42 will generate peak runoff at a later hour. After observing this precipitation, it is clear that hydrologic realizations in the probabilistic forecast calling for early peaks in runoff cannot be realized. The probability distribution is updated, eliminating these early peaks. The result is a significant reduction in the negative entropy of the forecast. The exponential rates of increase in the modal probability are combined with this one-time update of the range of the forecast distribution, producing a rate of change in entropy shown in Figure 14. This conditional distribution combined with the exponential rate of increased information content corresponding to certainty at hour 60 is used to simulate operation for the four storage allocations considered above.

As the example in section 4.3.1 suggests, a Bayesian update of the forecast distribution would continuously skew the forecast distribution. The change in the distribution depicted in Figure 14 would not be as abrupt. The single update at hour 43 reflects the observation of the maximum precipitation rate at hour 42, and the resulting certainty with which this observation eliminates the possibility of early peaks.

The effect of both exponentially increasing forecast skill and a one time truncation of the forecast distribution at hour 43 is quantified by the resulting maximum Luke discharge in simulated operation. Table 3 summarizes the operating simulations for each of the storage allocations considered. The results indicate the value of eliminating one form of hedging cost by improving the accuracy of the forecast. The maximum Luke discharge is dramatically reduced by real-time operations that use the truncated forecast distribution. For a fixed allocation of only 25,700 AF of flood control storage, the elimination of early peaking realizations combined with the forecast improvements, yields a flood peak of only 96,833 cfs. Although 10,300 AF of flood control storage has been reallocated, the combined forecast improvements lead to an increase in the maximum Luke discharge of only 499 cfs.

These operating improvements reflect the reduced hedging cost achieved through improved forecast skill. This reflects the added storage made available by eliminating the need to hedge against early peaking runoff hydrographs after hour 43. The stochastic operating rule has not been modified. Only the forecast information driving release decisions has changed. That is, real-time operation with improved forecasting can reduce the flood damage cost of a 10,300 AF storage reallocation by 7,230 cfs for the Bloomington SPF.

The robustness of the stochastic real-time operating rule is demonstrated in the following section. The stochastic operating rules are modified for a storage allocation in which part of the flood pool is explicitly shared with conservation storage. True multipurpose operation is achieved through non-structural hedging using real-time forecasts.

6. Jointly Allocated Storage

In this section the stochastic operating rules are modified to utilize jointly allocated storage. Non-structural hedging using real-time forecasts balances the risks of multipurpose operation. True multipurpose operation is achieved with real-time rules that operate a single storage allocation for multiple purposes.

When each hourly release decision is calculated, the only boundary conditions describing the system are the capacity of the flood control pool, the current volume of water in storage, and the past releases currently in the channel. Combined with the current hour's forecast, the operating rule determines the release that minimizes the expected maximum discharge downstream. In simulating a reallocation of reservoir storage, the only changes in the specification of the operating problem required are the reduction of both the capacity of the flood control pool and the initial storage. An alternate strategy for reallocating reservoir storage is developed in this section.

6.1 Balancing Risks from Joint Allocation

Reallocation studies are typically based on the permanent reallocation of flood control storage to another authorized project purpose. The effect of this reallocation is incorporated in the stochastic operating rule by reducing the initial volume of flood control storage and beginning flood control operations with the reallocated flood pool empty. The operating rule then prescribes releases for this smaller volume of flood control storage as reservoir inflow is observed and new forecast information is made available.

This suggests an alternate way to approach storage reallocations. Rather than taking storage away from flood control and allocating it to another project purpose, incremental volumes of both flood control and conservation storage could be explicitly allocated to shared operation. The initial conditions used to determine each hourly release in the stochastic operating rule guide the operation of jointly allocated reservoir storage. This jointly allocated volume could be emptied for flood control when needed and operated to refill on the drawdown phase of flood control operations.

The stochastic operating rule provides a natural tool to operate this shared storage in real-time based on forecast information. Consider, for example, a reallocation of 5600 AF of flood control storage to a joint flood control-conservation pool. When 5,600 AF of storage is reallocated to this new, shared use, the capacity of the flood control reservoir remains unchanged at 36,000 AF. The shared reallocation imposes an initial storage condition that is greater than zero. This reallocation would impose an initial storage of 5,600 AF for the 36,000 AF flood pool. Guided by real-time forecasts of runoff, the operating rule may store inflow, release inflow, or release shared storage, as conditions change.

In order to make the shared storage available for conservation purposes, flood control operations are constrained to refill to the initial storage level at the end of the operating horizon. When flood control storage is being evacuated, refilling jointly allocated storage is a non-restrictive requirement.

The risk in refilling jointly allocated storage is that insufficient inflow is stored to make this possible. Making refill a binding constraint means that the lowest forecasted inflow determines the volume of shared storage that could be released for flood control purposes, while allowing shared storage to refill. If there is a possibility that there will be no high runoff event (for example when a storm system poses a threat to the region but has not yet affected the entire basin) no drawdown of shared storage will occur.

Real-time operation of jointly allocated storage supports the need for more conservation storage, while mitigating the loss of flood protection such a reallocation would typically produce. In contrast to operations in which the allocation for each authorized project purpose is operated independently, the shared storage allocation leads to true multipurpose reservoir operation. Extra conservation storage provides an added level of reliability for water supply or flow augmentation benefits; the opportunity to temporarily evacuate shared storage provides added reliability for flood control operation. This added reliability for both conservation and flood control purposes is provided by the same volume of jointly allocated storage. By using the same volume of storage to provide marginal reliability for several purposes, the benefits for each purpose can be increased without requiring any additional reservoir storage. The key to successful joint storage allocation is reservoir operating rules that can fully utilize real-time forecast information.

Successful dynamic allocation results from operating decisions that explicitly balance the risks from multipurpose operation. The simplicity with which dynamic allocation can be achieved demonstrates the robustness of the real-time operating rule. The boundary conditions that define dynamic allocation are entirely analogous to those determining the release during any hour of flood control operation. Dynamic allocation represents a special case of the general real-time flood control operating problem.

The fundamental difference between the operation of a dynamic allocation and a static allocation of flood control storage is the evacuation of jointly allocated storage early in the flood event. In contrast to all the operating simulations reported in section 5, dynamic allocation results in releases on the rising limb of the flood hydrograph that increase the downstream discharge above natural, unregulated levels. When these pre-releases exceed flood stage, releases prescribed under dynamic allocation result in earlier downstream flood damages than natural, unregulated flow conditions. In return for making these pre-releases, additional storage is made available later in the storm, reducing the maximum downstream discharge to levels that would otherwise be unattainable.

The pre-release decision is a typical hedging problem. In this case the hedging decision is to raise downstream stages with certainty, in order to reduce the expected value of future damages. The value of making pre-releases in flood control operation has long been recognized. The risks associated with pre-release strategies have usually been judged unacceptably high in most multipurpose reservoirs. Due to the uncertainty of the runoff forecast early in a flood event, the value of the additional storage from pre-release strategies cannot be quantified with certainty. The potential damages that could be caused by pre-releases place added responsibility on the reservoir operator. If damages were caused by pre-releases but the additional storage that resulted was not used, the consequences would be unacceptable. An alternate pre-release strategy would drawdown conservation storage at a low enough rate to ensure that no damages would be caused downstream. This safe drawdown rate would require drawdown over a period of days or weeks. If the anticipated flood flows were not realized, the conservation pool might not refill, raising the risk of failure for conservation purposes to unacceptably high levels.

The risks from both incurring unnecessary damages, as well as failure of conservation operation, arise largely from uncertainty in the runoff forecast. These risks have usually been viewed as unacceptably high. Where runoff can be forecasted with dependable skill, pre-release strategies are less risky. On the Kansas river where flood peaks precede the peak on the Mississippi at Cairo, pre-releases are made within storm events (Coomes, 1969). For snowmelt floods in which the precipitation that will generate runoff can be accurately estimated from snow pack surveys, great potential exists for dynamic allocation and pre-release strategies. Indeed seasonal rule curves in regions dominated by snow-melt floods represent a form of dynamic allocation. The stochastic operating rules demonstrated in this section extend pre-release strategies to real-time operation. The extent to which jointly allocated storage will be evacuated early in a flood event is determined by the solution of a real-time hedging problem, using probabilistic forecast information. For a joint storage allocation, the stochastic operating rule uses real-time forecasts to balance the risks associated with pre-release strategies, against the value of using jointly allocated storage for flood control.

The stochastic operating rules are driven not by the expected value, but by the probability distribution of runoff. The probability distribution of runoff defines the operating risks associated with any release. The pre-release resulting from dynamic allocation is driven by the real-time forecasts of reservoir inflow and uncontrolled runoff. Pre-releases will only be made when the probable damage reduction resulting from the additional storage that is made available, is greater than the cost of the increased downstream discharge required to evacuate jointly allocated storage. The operating rule explicitly balances the risk of causing unnaturally high stages downstream against the probable increase in flood protection. The risk of failing to refill jointly allocated storage is similarly managed through the use of the runoff forecast. One of the feasibility conditions imposed in the stochastic optimization problem used to determine the current period's release is the refill of jointly allocated storage at the end of the operating horizon. The runoff forecast that determines the feasibility of refill is the minimum runoff volume over the operating horizon. In other words, the maximum drawdown of jointly allocated storage will be bounded by the minimum forecasted runoff volume, ensuring the feasibility of refill.

6.2 Simulated operation of jointly allocated storage

The use of real-time forecasts and the stochastic operating rules with jointly allocated storage is demonstrated for the Bloomington SPF. The usable capacity of the flood pool is maintained at 36,000 AF for all reallocations. Only the fraction of the pool that is initially full is modified. Dynamic allocations of 5,600 AF, 10,400 AF, and 12,200 AF are simulated using both the exponential reduction in negative entropy, and the truncation of the runoff parameter distribution, discussed in section 5.2.3 The increase in the probability of the mode of the runoff forecast grows exponentially from hour 40 to hour 60. In addition, at hour 43 the early peaking realizations are judged infeasible and are eliminated from the joint parameter distribution, as in section 5.2.3 above. The operating results from a combination of dynamic allocation and improving forecast skill is summarized in Table 4.

The results summarized in Table 4 illustrate not only the value of dynamic allocation, but also the way in which flood control and conservation benefits are balanced in dynamic allocation. Compared to a static reallocation with forecast improvement, the dynamic allocation reduces the downstream flood peak for all joint storage allocations. No dynamic allocation is able to achieve the level of flood peak reduction that results from the combined forecast improvements operated for the original 36,000 AF allocation. However, the joint allocations dependably provide additional conservation storage that would otherwise be unavailable.

Although the jointly allocated storage is partially utilized in operating each dynamic allocation, caution is exercised in the early drawdown of jointly allocated storage. The overall volume of empty flood control storage approaches, but never equals the storage available under the current 36,000 AF allocation. For this reason, flood peak reduction from a joint allocation will approach but never equal the protection provided by the full 36,000 AF of flood control storage.

As the jointly allocated volume is increased, the downstream flood peak is expected to similarly increase. The increase in the downstream flood peak is not, however monotonic with increasing volume of jointly allocated storage. The downstream flood peak increases from 85,244 cfs to 92,005 cfs when 5,600 acre feet of original flood control storage are reallocated to jointly allocated storage. The joint allocation of 5600 AF provides a small reduction in the downstream flood peak compared to a fixed reallocation of 5600 AF of flood control storage. When the joint storage allocation is increased to 10,400 AF, the downstream flood peak is reduced to 87,630, nearly equal to the flood peak reduction achieved with the original 36,000 AF flood pool. This curious result reflects the way in which the operating rule uses real-time forecasts to balance the risks of operating dynamically allocated storage.

For the reallocation of 5,600 AF to joint storage, the incremental reduction in the flood peak from drawdown of joint storage is small compared to the increase in the downstream stage these releases would cause. The releases prescribed early in the flood event barely drawdown jointly allocated storage. This results in the availability of a small volume of incremental storage later in the flood. The result is an insignificant increase in level of flood protection achieved from a joint allocation of 5600 AF.

For a joint allocation of 10,400 AF, the expected reduction in the flood peak becomes significant and the early releases during the SPF are much larger than reservoir inflow, evacuating most of the jointly allocated storage. These increased releases have the additional effect of raising the downstream boundary condition in the channel during later operation. The rise in stage effectively reduces the need to hedge against low magnitude peaks. Protection against relatively small flood peaks is abandoned to ensure adequate storage is available to reduce large peaks. The result is a maximum discharge at Luke that is only 2000 cfs higher than the peak achievable with 36,000 AF of storage.

Increasing the joint allocation to one third of the original flood pool results in a downstream flood peak of 90,050 cfs. This increase reflects the inability of the operating rules to fully pre-release the jointly allocated storage volume. Although the flood peak from the 12,200 AF joint allocation has increased, the 90,050 cfs Luke discharge is less than any regulated flood peak resulting from a fixed allocation, operated with the combined forecasts, as in Table 3. Compared to the original SPF operation with a 36,000 AF flood pool, the operation of a 12,200 AF joint allocation provides a 6,284 cfs reduction in the flood peak as well as an extra 12,200 AF of conservation storage.

7. CONCLUSION

This report examined the potential for improving the operation of flood control reservoirs through the combined use of stochastic real-time reservoir operating rules and real-time forecasting. Where operating gains can be realized in flood control operation, the reliability of flood protection can be increased, flood damages can be reduced, and the volume of storage required to provide flood protection can be reduced, freeing storage in existing reservoirs for alternate uses.

The Potomac River is the second largest source of freshwater to the Chesapeake Bay, after the Susquehanna River. Additional reservoir storage on the Potomac River could provide a supplemental source of freshwater for the Chesapeake Bay during low flow periods in the summer and late fall. If flood control storage on the Potomac could be operated more efficiently, reservoir storage in existing impoundments could be operated to supplement freshwater flows without sacrificing any flood protection.

Jennings Randolph Reservoir on the North Branch of the Potomac River is the largest impoundment in the Potomac River Basin. The value of forecasting skill in conjunction with real-time stochastic operating rules was demonstrated through simulated flood control operation of the 36,000 acre-foot flood control pool in Randolph Reservoir. The operating simulations demonstrate that the combination of improved forecasting and real-time operating rules can improve reservoir operation for flood control in three ways:

- Reservoir releases must hedge against uncertain hydrologic conditions during a flood. Operating rules that incorporate the accuracy of the runoff forecast in release decisions will deliver increased protection from an improvement in forecast accuracy.

- The operating forecast represents a probability density function of runoff. Operating rules that explicitly use this density function in determining releases reduce the cost of risk taking through non-structural hedging.

- Releases that explicitly balance operating risks in real-time can achieve true multipurpose reservoir operation through the operation of jointly allocated reservoir storage.

This report considered the reallocation of 5,600, 10,300, an 12,200 AF of Randolph flood control storage, to conservation purposes. Four levels of forecast skill were simulated and used to compare the operating consequences of both a fixed allocation of flood control storage, as well as a joint allocation in which a portion of the flood pool was allocated to both flood control and conservation purposes. Figure 15 shows the resulting trade-off of flood damages, reported as maximum discharge at Luke,Md., and storage reallocation, for all operating simulations. From the operating simulations, the following conclusions can be drawn:

- The reallocation of a portion of the existing flood control pool in Randolph reservoir to conservation purposes will increase the downstream flood damages. The use of improved forecasts, in conjunction with operating rules that can respond to these forecast improvements can reduce, and in some cases eliminate the loss of flood protection resulting from a storage reallocation

- For any volume of flood control storage considered in this report, operations with forecasts that are updated during the flood event significantly reduces flood damages compared to operating strategies with constant hedging

- The forecast improvements that yield the greatest marginal reductions in flood damages, are also the improvements that are most easily achieved. Additional incremental improvements in forecast skill will be increasingly difficult to achieve, and will have small but positive incremental value.

- Operation of jointly allocated storage significantly improves the benefits that can be achieved through the use of stochastic real-time operating rules and real-time forecasting. For the operating simulations considered in this report, maximum forecast skill combined with the joint allocation of 12,200 AF of flood control storage to jointly allocated storage could reduce the Luke flood peak by over 6,000 cfs for the bloomington SPF while providing an additional 12,200 AF of conservation storage to enhance the flow of freshwater to the Chesapeake Bay.

Figure 16 shows the unit reduction in the flood peak at Luke,Md. for all of the operating simulations considered in this report. Comparing the operating simulations indicates:

- o Marginal reallocations of flood control storage will increase the unit protection provided by the remaining storage. The incremental flood protection provided by the last increment of storage in a flood control allocation is storing peak damaging inflows with the longest duration. As incrementally damaging inflows are released due to a loss of storage, the duration of the flows that are stored decreases, making the remaining storage more productive per acre-foot while raising the downstream flood peak.

- o Operation with stochastic real-time operating rules and improving real-time forecasts dramatically increases the flood protection offered per acre-foot for all storage allocations. The increased productivity of flood control storage achieved with real-time operation and forecasting can more than offset the loss of flood protection resulting from a reallocation operated with constant hedging.

These results indicate that great potential exists for utilizing real-time forecasting and operating procedures to support additional uses of existing flood control storage. Existing cost-effective technology can provide significant skill in estimating the probability density function of runoff during a flood. Real-time operating rules that can use probabilistic forecast information are a powerful tool for improving the utilization of existing reservoir storage. On the Potomac river up to 12,000 AF of reservoir storage could be made available to enhance the flow of freshwater to the Chesapeake Bay without compromising existing reservoir operation.

REFERENCES

Bras, R.L. and I. Rodriguez-Iturbe (1985) Random Functions and Hydrology. Addison-Wesley Publishing Company.

Coomes,R.T. Flood Regulation of Kansas River Basin Reservoirs. In Proceedings of a Seminar on Reservoir Systems Analysis. U.S. Army Corps of Engineers, Hydrologic Research Center. Davis, California. November, 1969

Eagleson, P.M (1967) Optimum Density of Rainfall Networks. Water Resources Research, 3(4) pp 1021-1034

Grace, R.A. and P.S. Eagleson (1966). The Synthesis of Short-Time Increment Rainfall Sequences. MIT Tech. Report #91. MIT, Cambridge, Massachusetts

Grayman W.M. and P.M. Eagleson (1971). Evaluation of Radar and Raingage systems for Flood forecasting. Ralph M. Parsons Laboratory Technical Report 138, MIT.

Hendrich, R.L. and G.H. Comer (1970) Space Variations of Precipitation and Implications for Raingage Network Design. Journal of Hydrology vol. 10, pp151-163

Hogan,K. (1987) Application of facility Location Models to Raingage Network Design. Paper Presented at the Regional Science Association Meeting, Baltimore, Maryland, November 1987.

Huff, F.A. (1967) Time Distribution of Rainfall in Heavy Storms. Water Resources Research vol 3 no. 4 pp 1007-1019

Pilgrim, D.H., I. Cordery, and R. French (1969) Temporal Patterns of Design Rainfall for Sydney. Civil Engineering Transactions, Institution of Engineers

Schwartz, S.S. and D.P. Sheer (1984) Design and Operating Rules for reservoir storage to offset Future consumptive Use. Proceedings of the ASCE WRPM Div. Specialty Conference on computers in Water Resources, Buffalo,N.Y.

Smith, J.A.,S.S. Schwartz, J.W. Whitcomb, and J.B Boland (1987) Multiobjective Reservoir Operation Using Forecasts of Water Supply and Water Use. USGS Water Resources Research Grant No. 14-08-0001-g1145

Tsonis, A.A. (1987) Determining Rainfall Intensity and Type from GOES Imagery in the Midlatitudes. Remote Sensing of the Environment, 21:29-36

Zawadzki, I.I. (1973) Errors and Fluctuation of Raingage estimates of Areal Rainfall. Journal of Hydrology, no. 18, pp243-255.

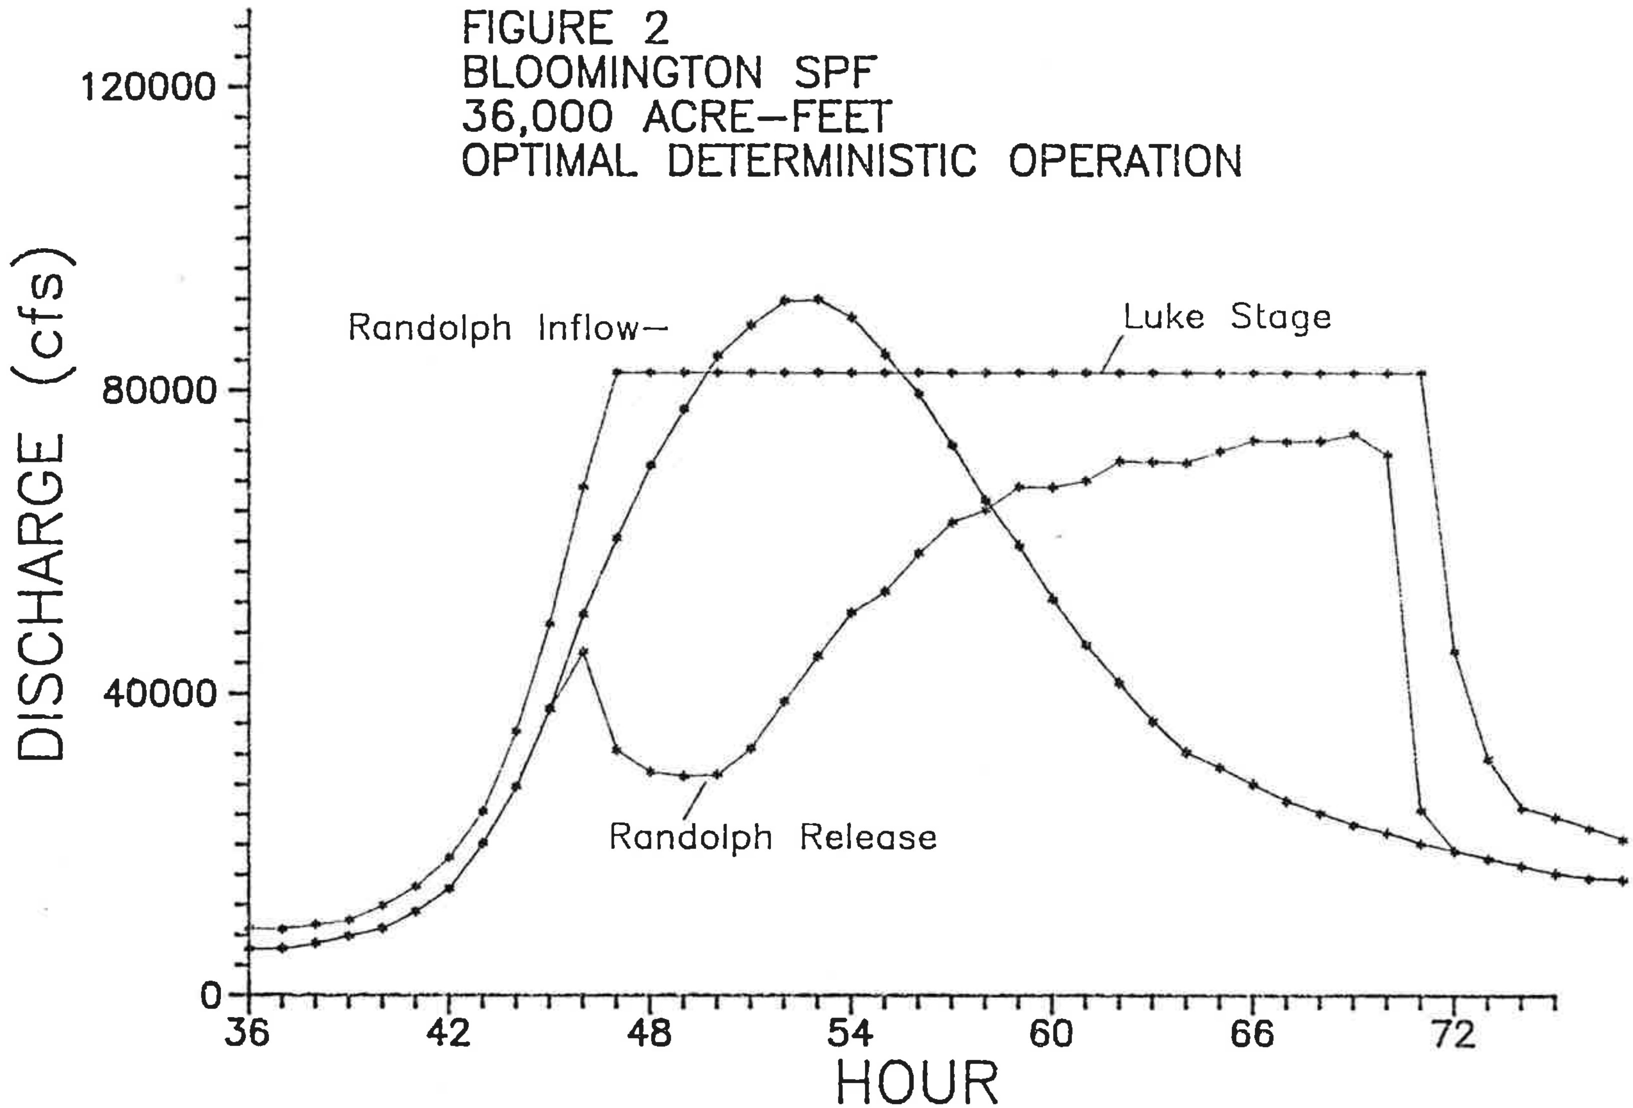
FIGURE 2
BLOOMINGTON SPF
36,000 ACRE-FEET
OPTIMAL DETERMINISTIC OPERATION
Randolph Inflow–
Luke Stage
Randolph Release
DISCHARGE (cfs)
120000
80000
40000
0
36
42
48
54
60
66
72
HOUR

FIGURE 3
BLOOMINGTON SPF
36,000 ACRE–FEET
DETERMINISTIC AND STOCHASTIC OPERATION

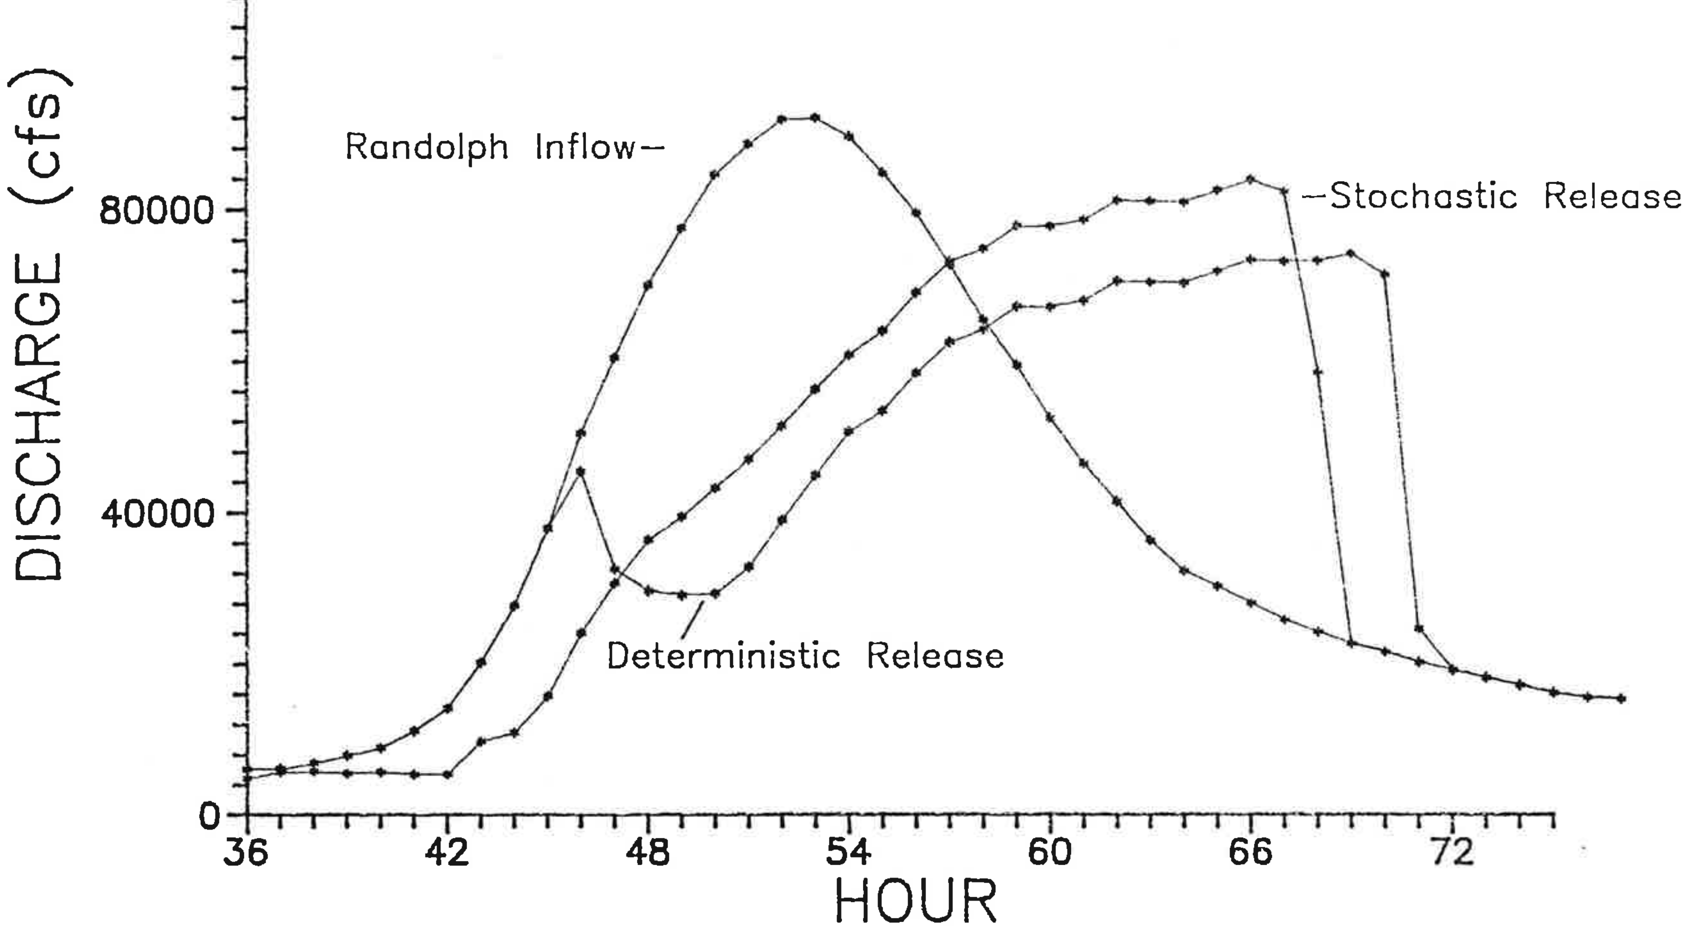

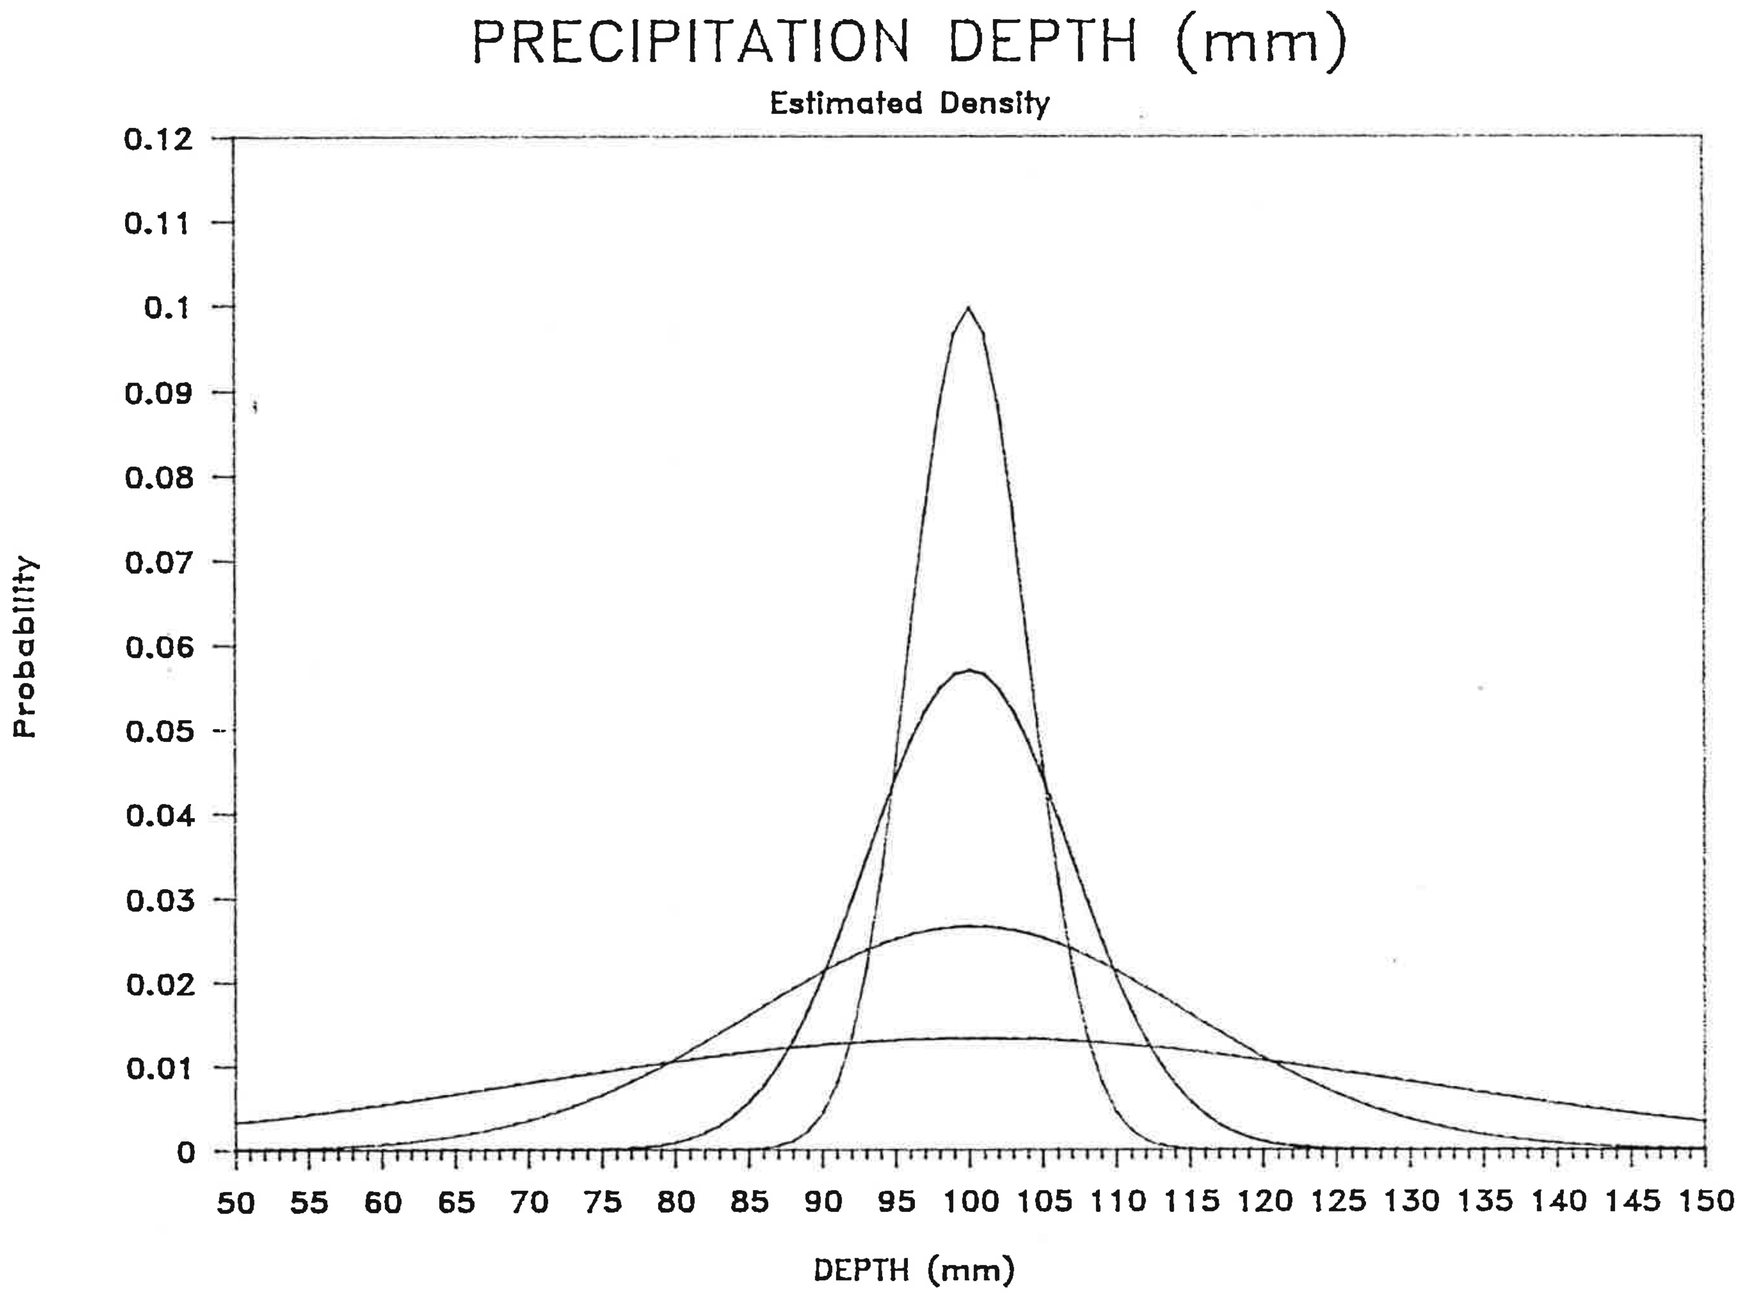
PRECIPITATION DEPTH (mm)
Estimated Density
Probability
0.12
0.11
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
50 55 60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150
DEPTH (mm)

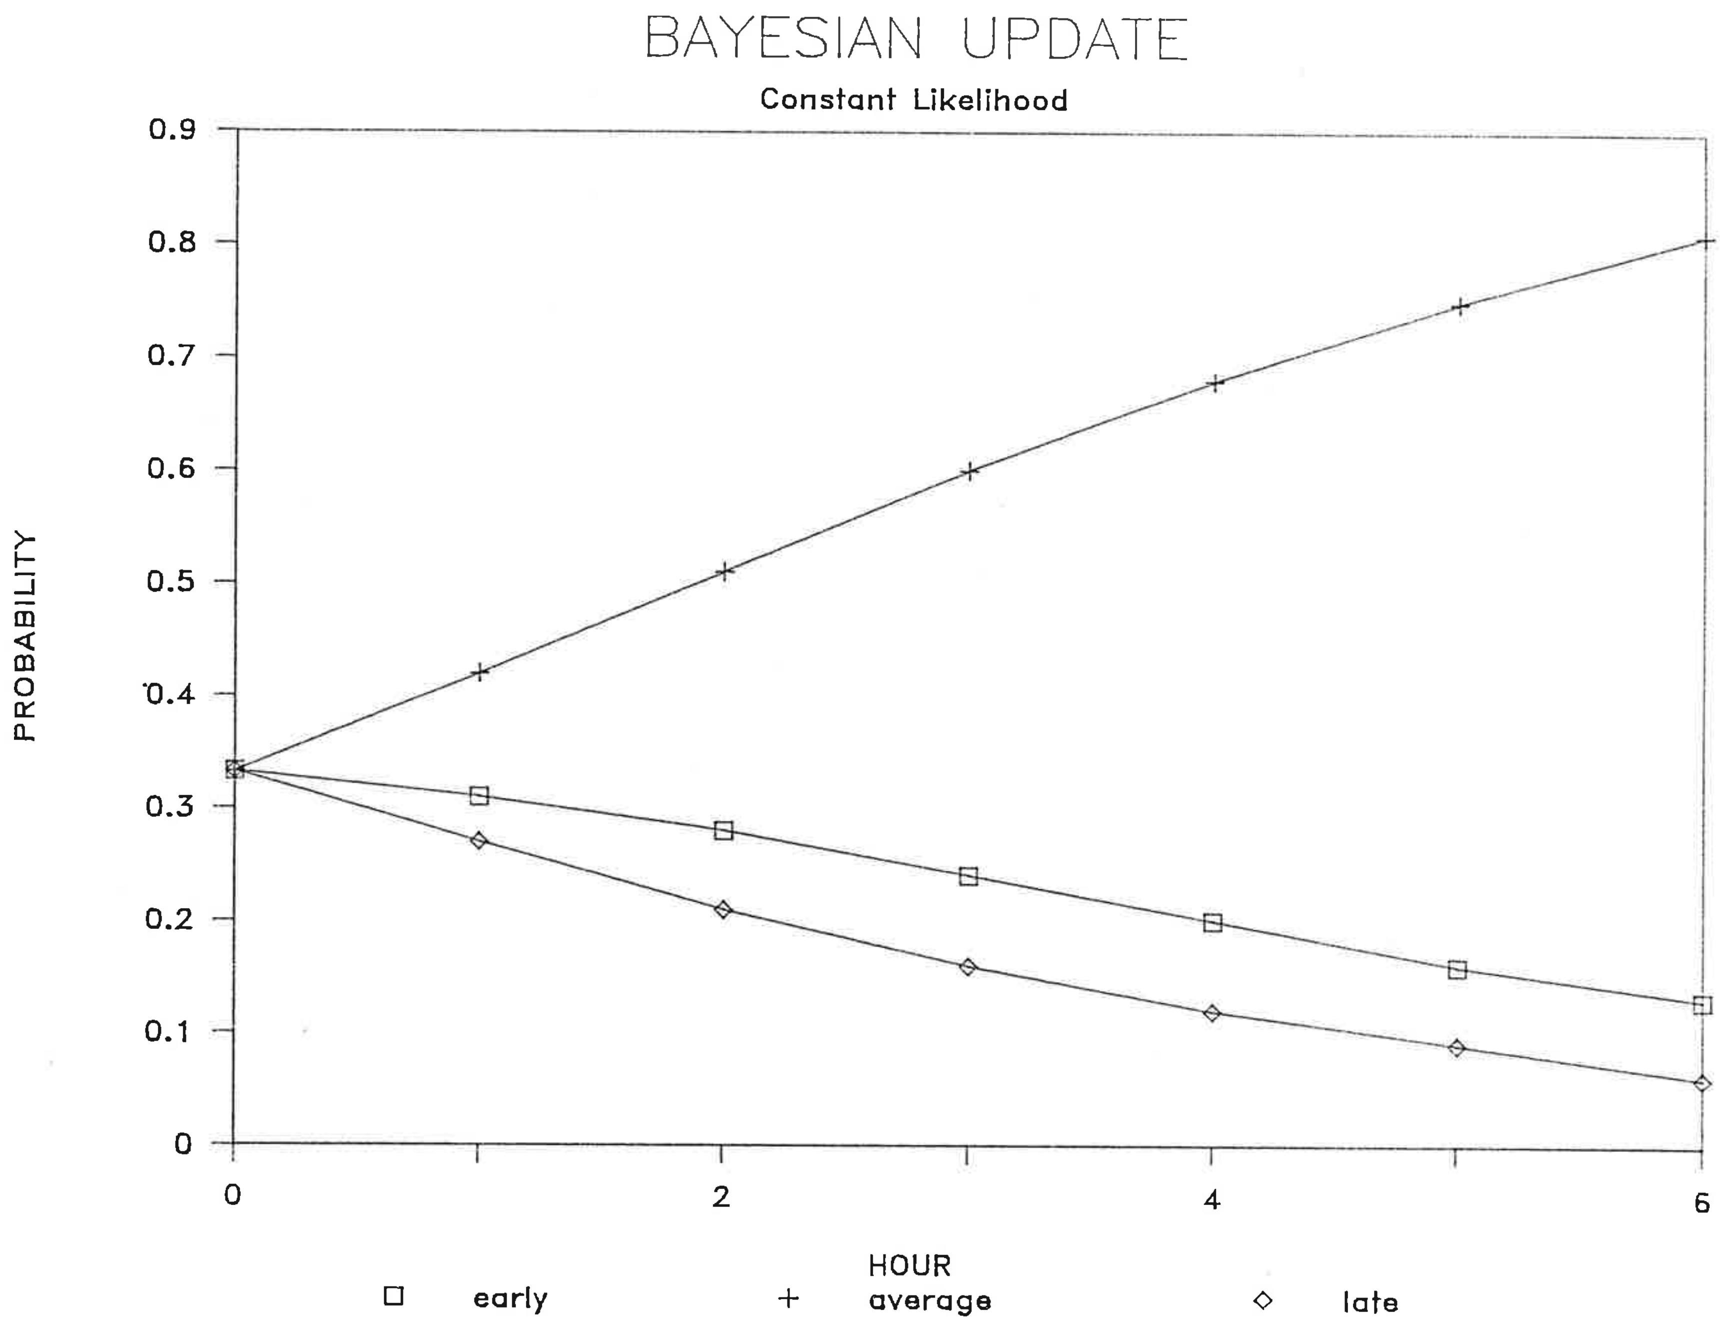
BAYESIAN UPDATE
Constant Likelihood
PROBABILITY
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
2
4
6
HOUR
early
average
late

BAYESIAN ESTIMATION OF PEAK TIME

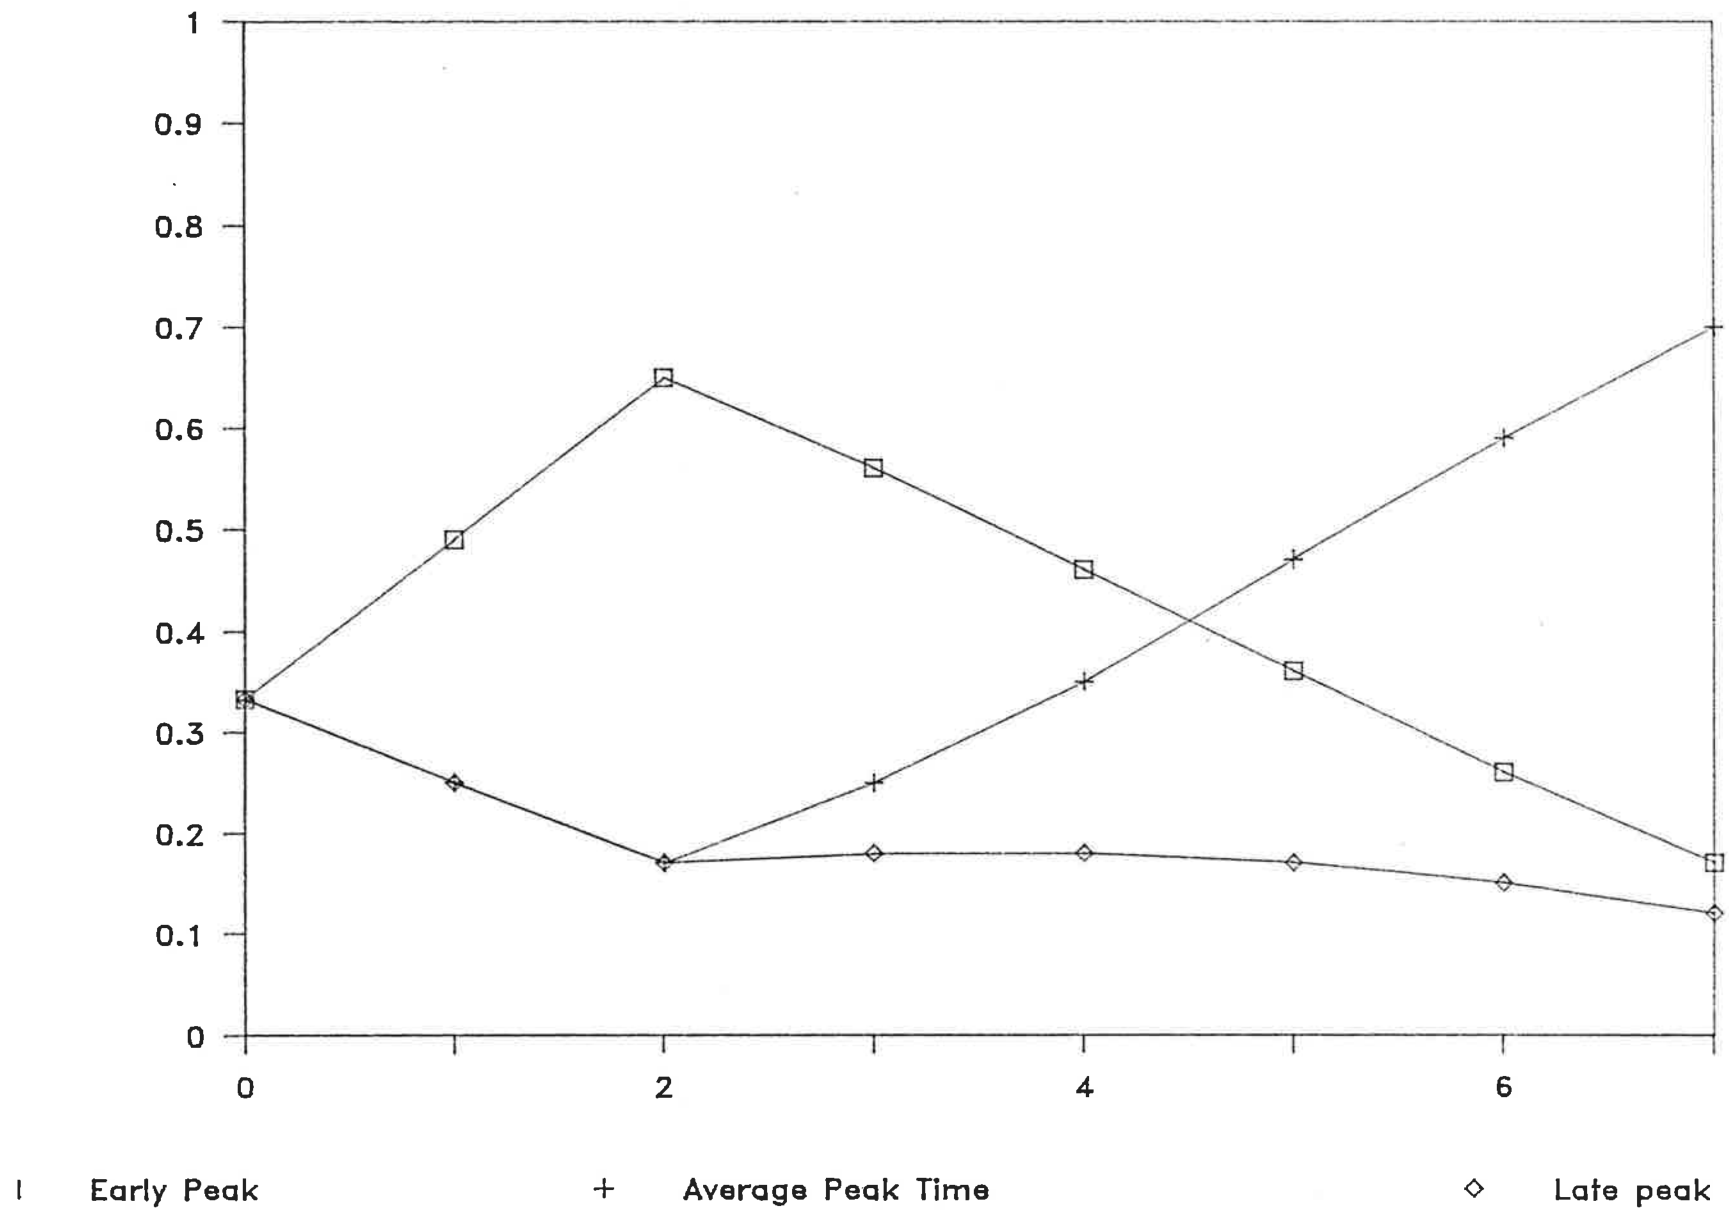

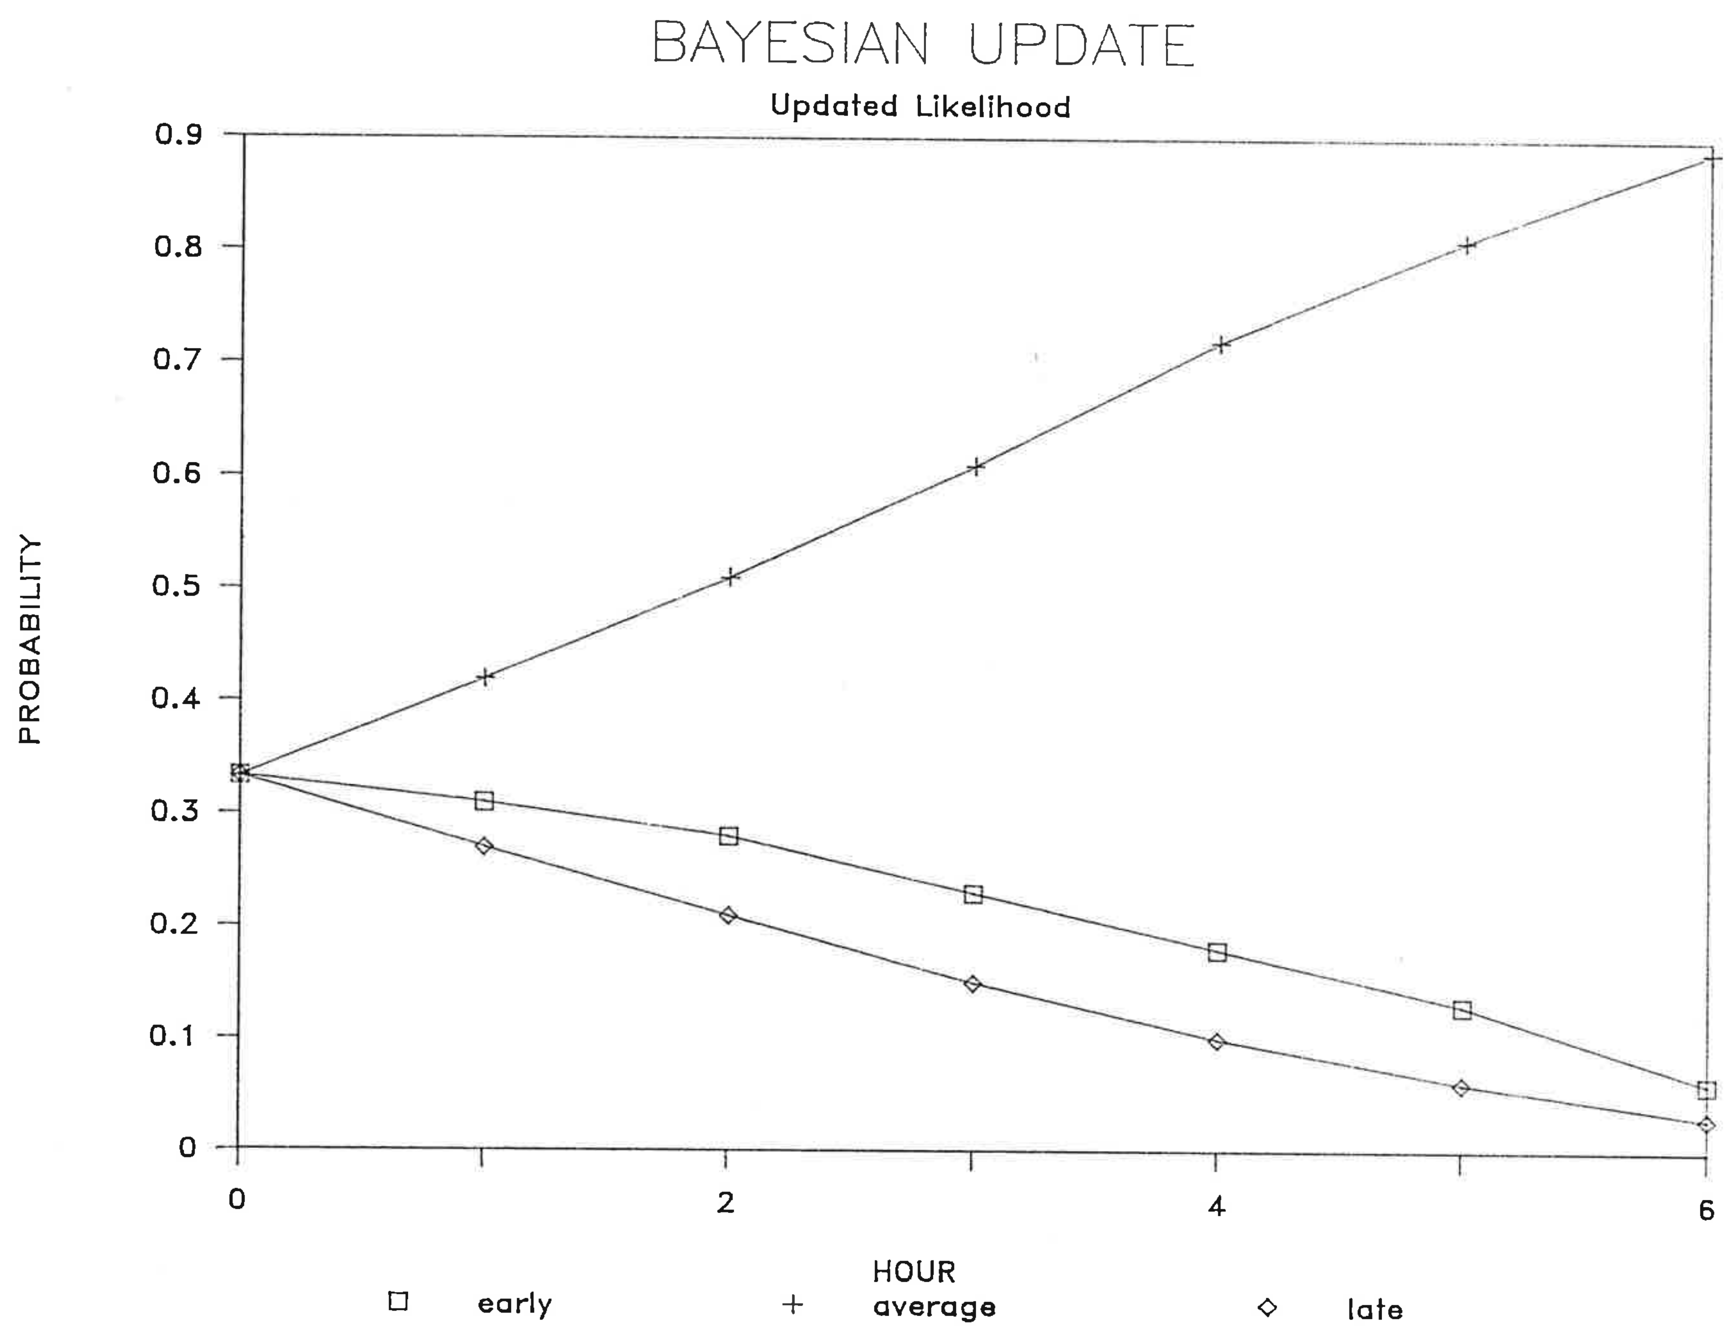

BAYESIAN UPDATE
Updated Likelihood
PROBABILITY
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
2
4
6
HOUR
early
average
late

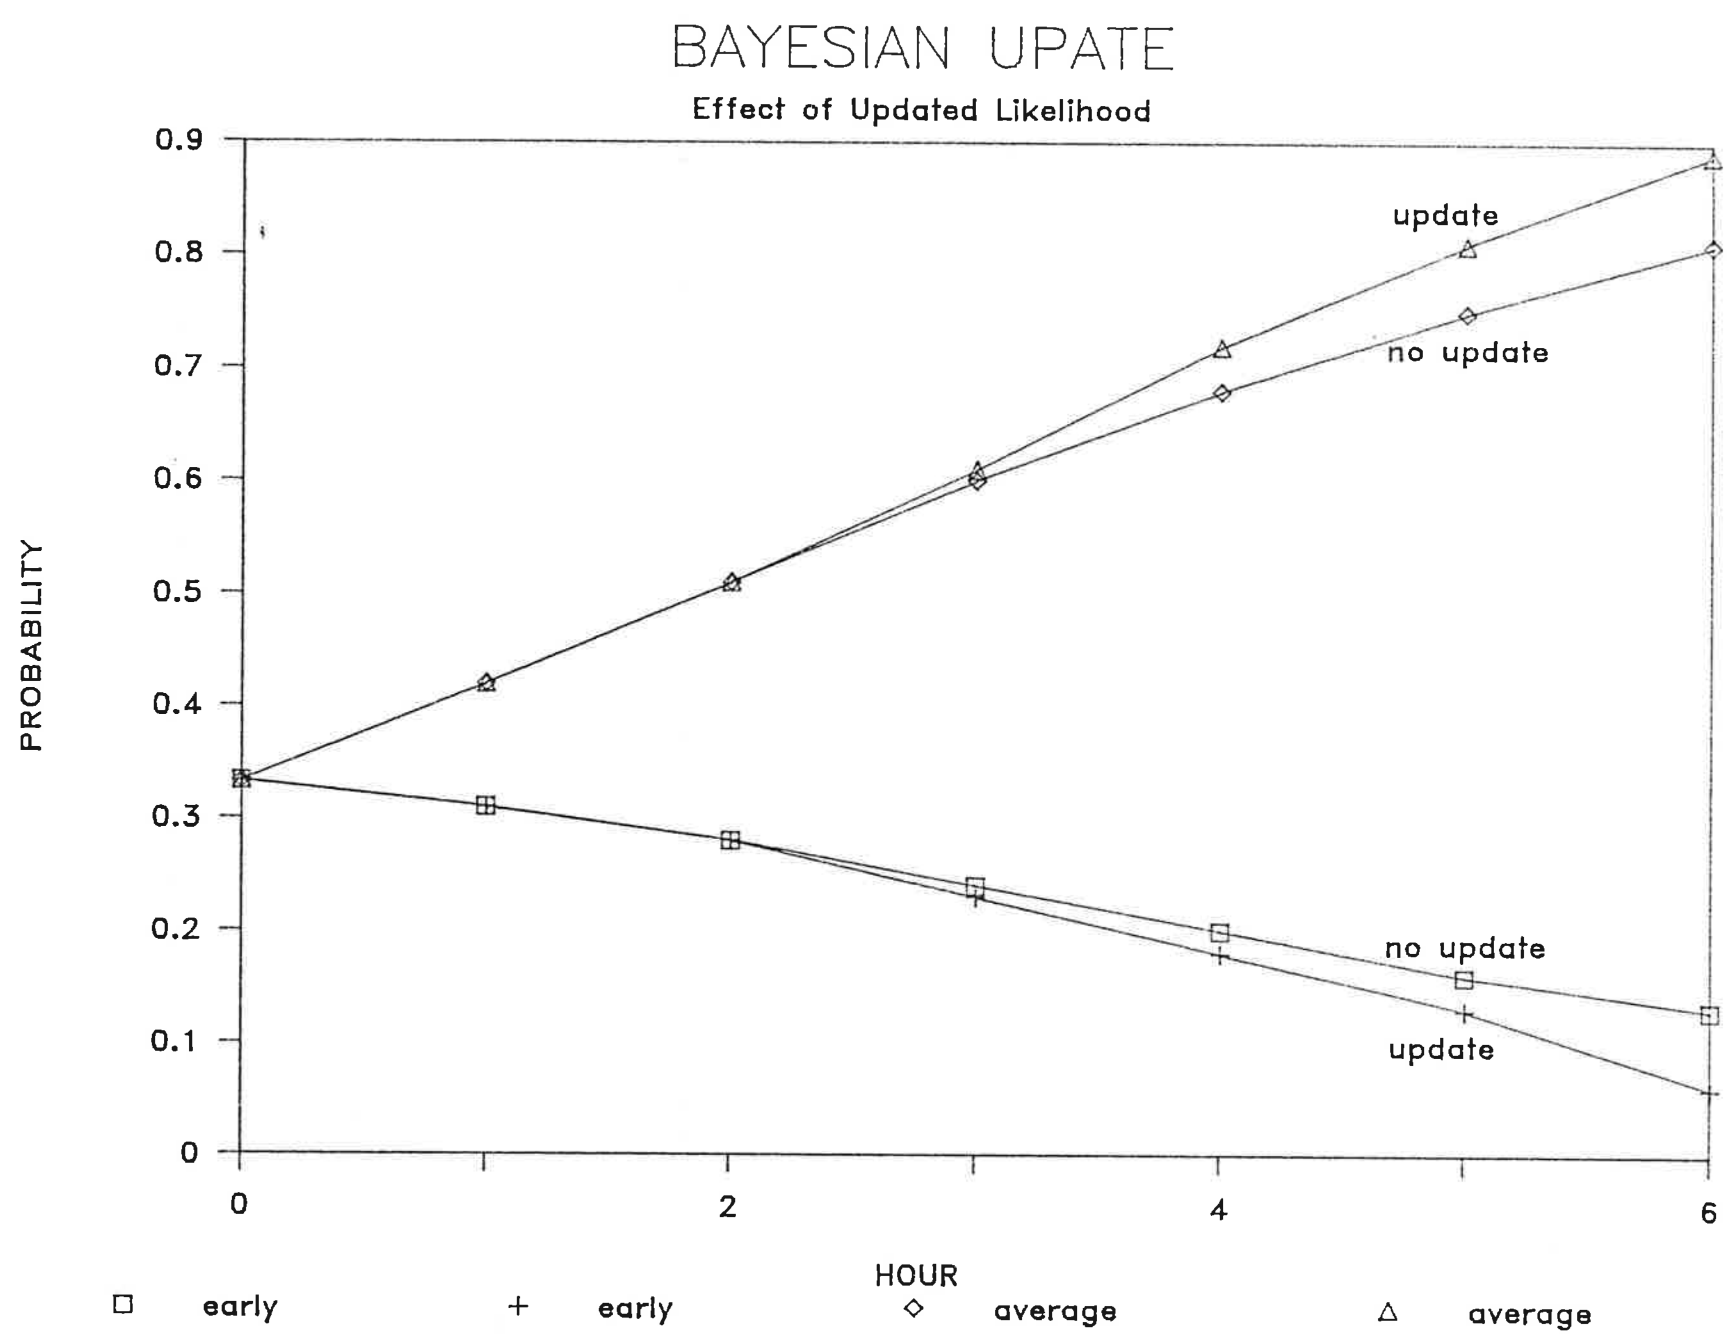
BAYESIAN UPATE
Effect of Updated Likelihood
PROBABILITY
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
update
no update
no update
update
0
2
4
6
HOUR
□ early
+ early
◇ average
△ average

PRECIPITATION DEPTH (MM)

Conditional Density

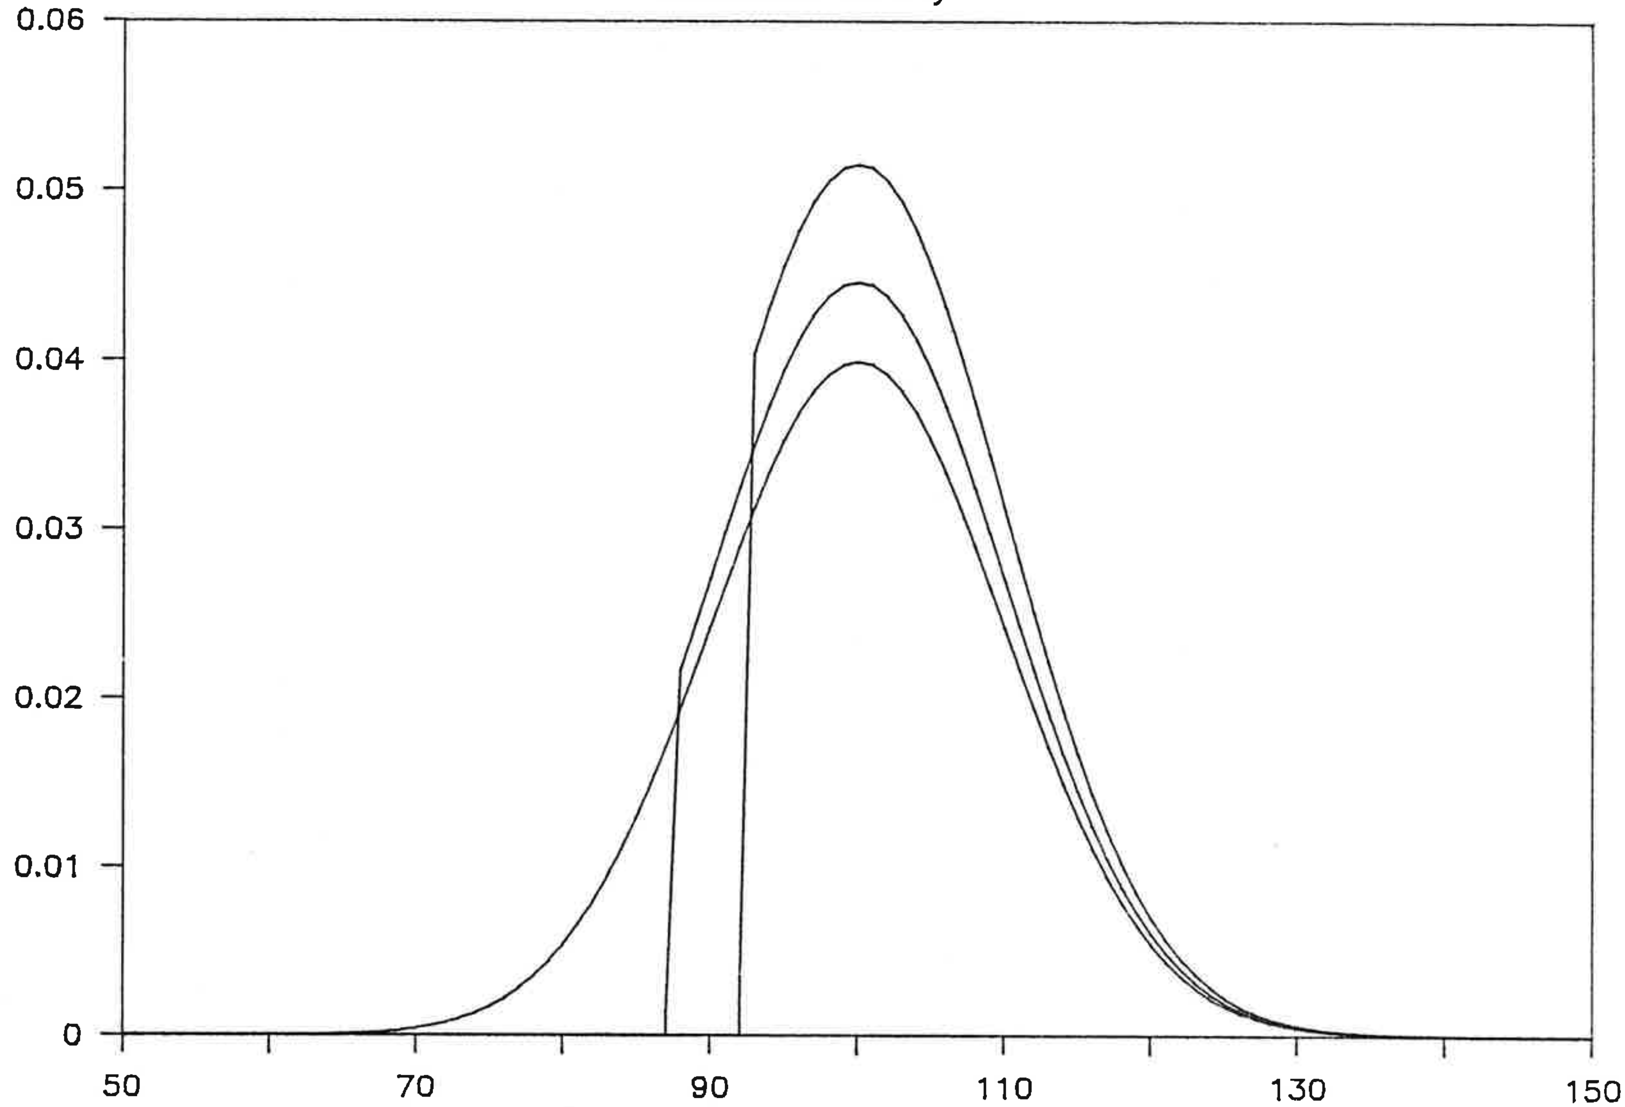

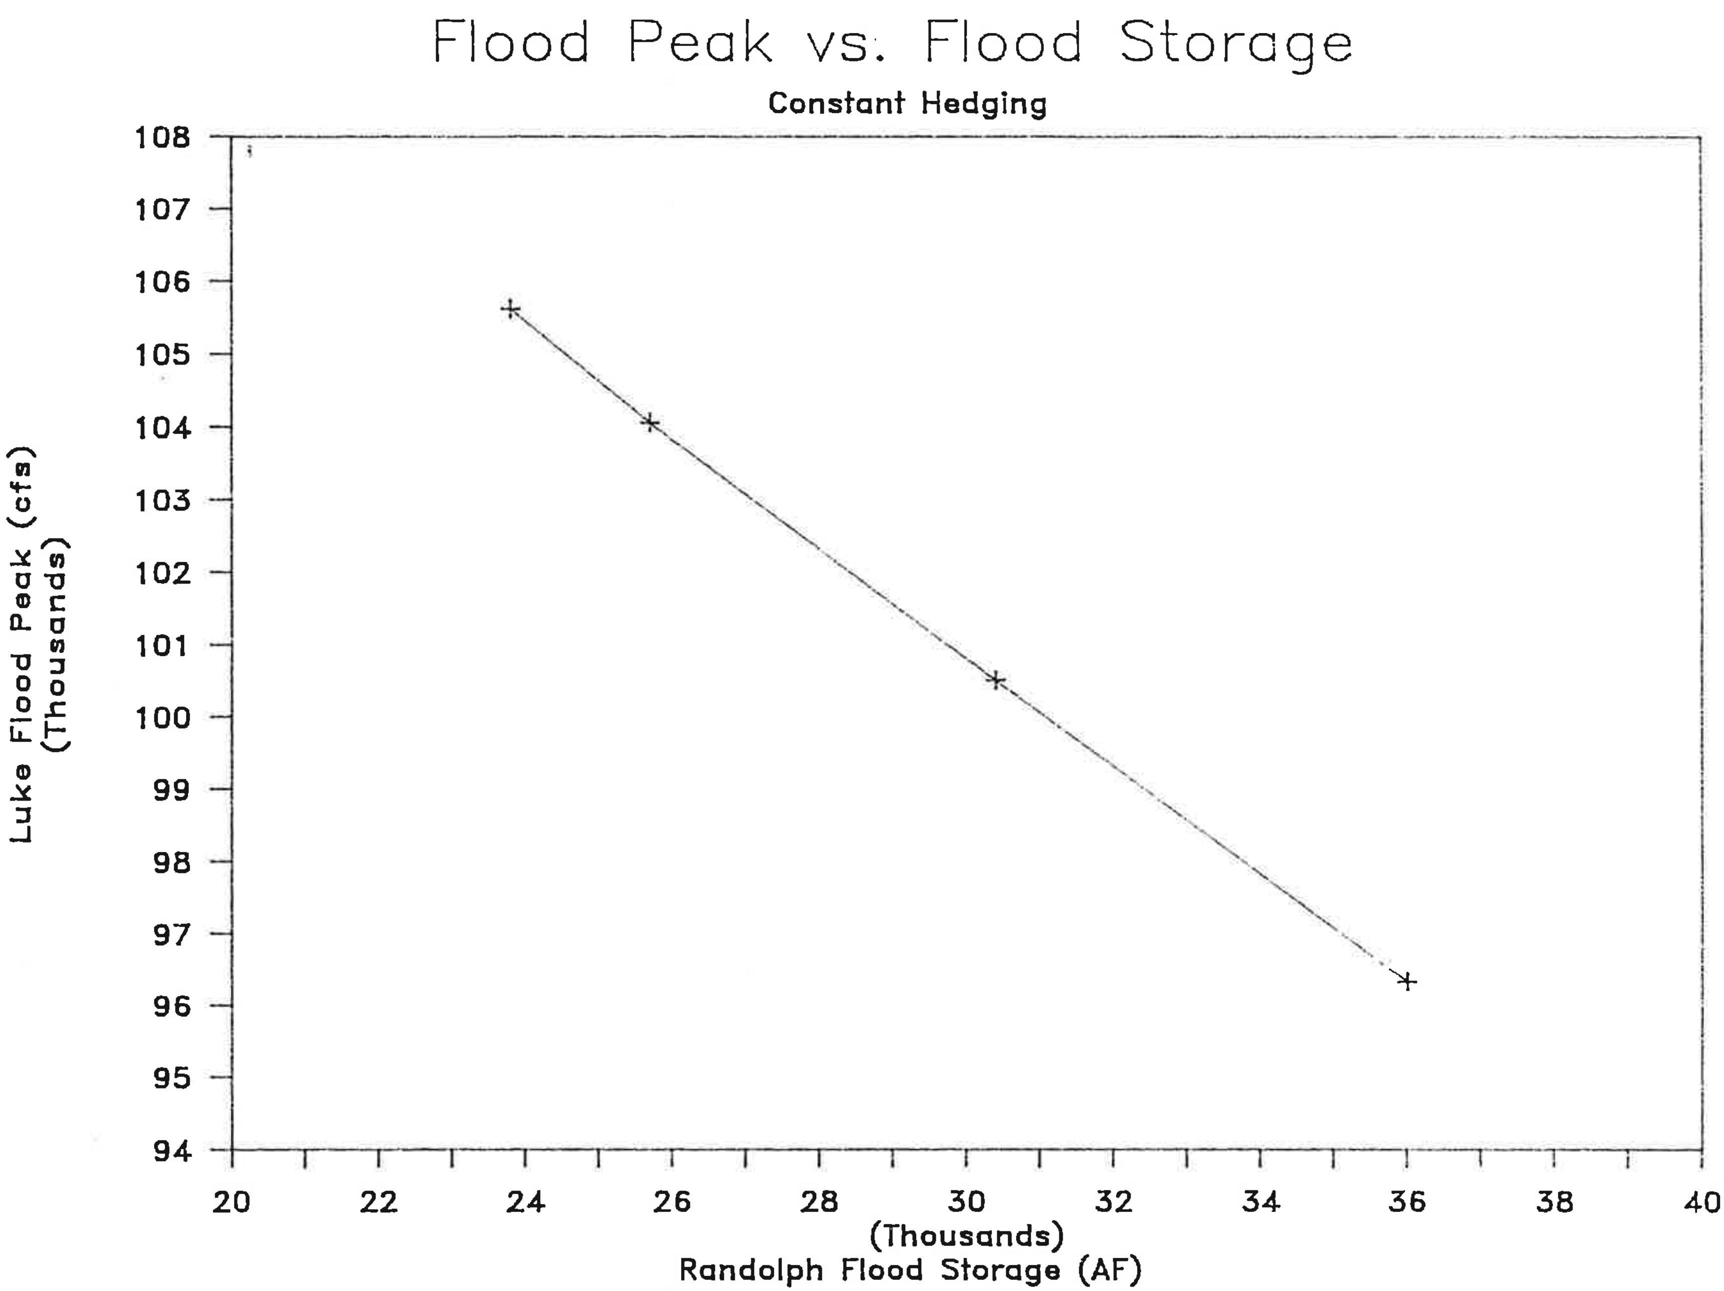
Flood Peak vs. Flood Storage
Constant Hedging
Luke Flood Peak (cfs)
(Thousands)
108
107
106
105
104
103
102
101
100
99
98
97
96
95
94
20
22
24
26
28
30
32
34
36
38
40
(Thousands)
Randolph Flood Storage (AF)

STORM HYETOGRAPH

Bloomington SPF

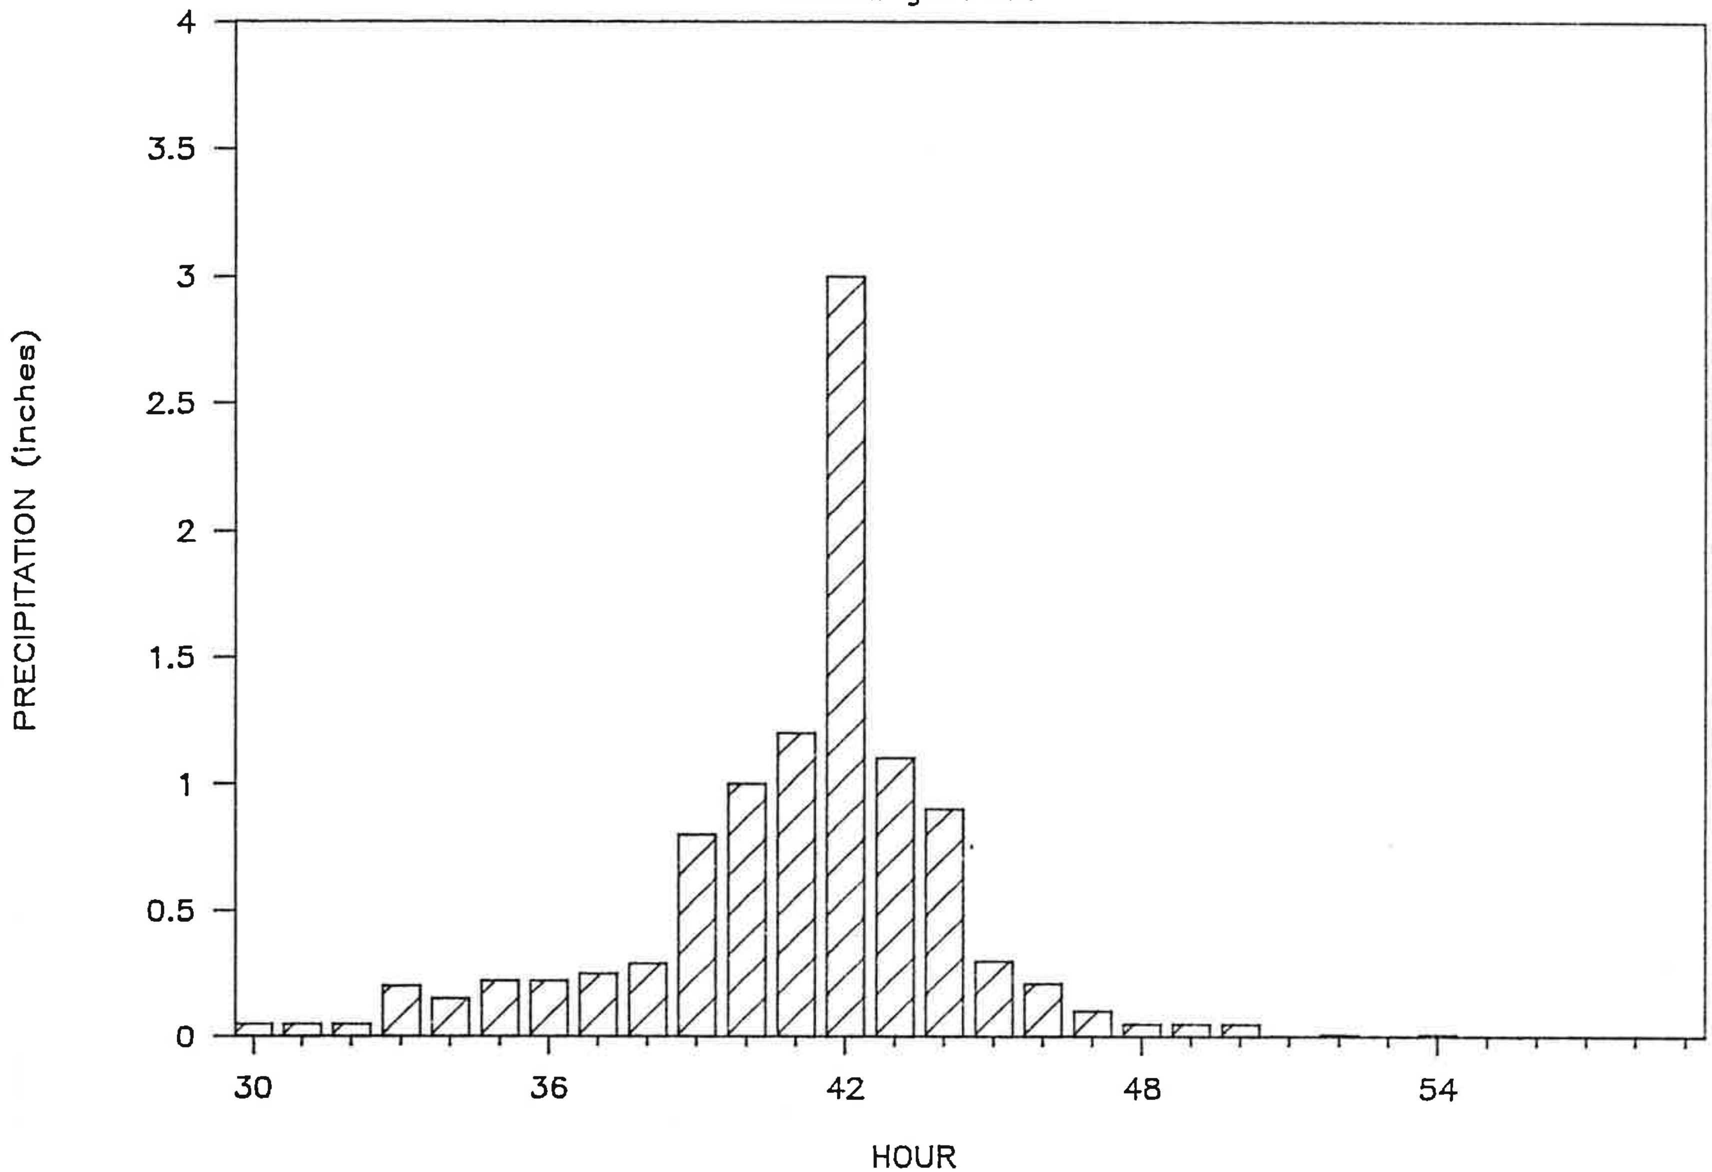

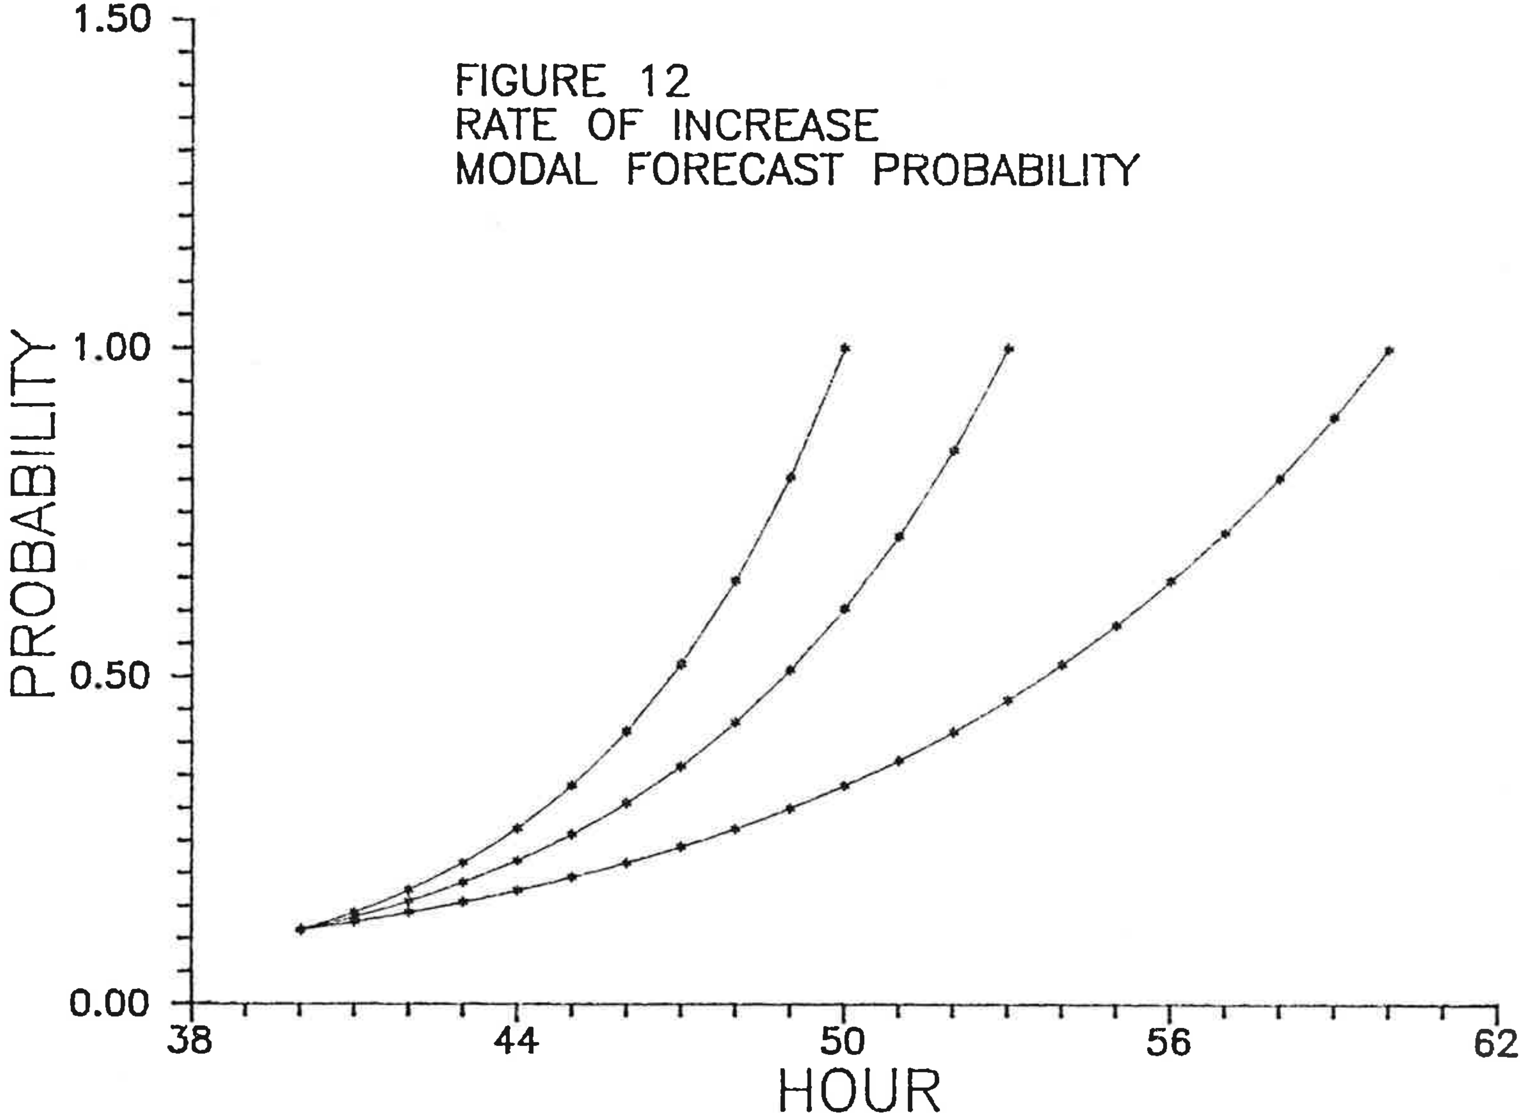

FIGURE 12
RATE OF INCREASE
MODAL FORECAST PROBABILITY

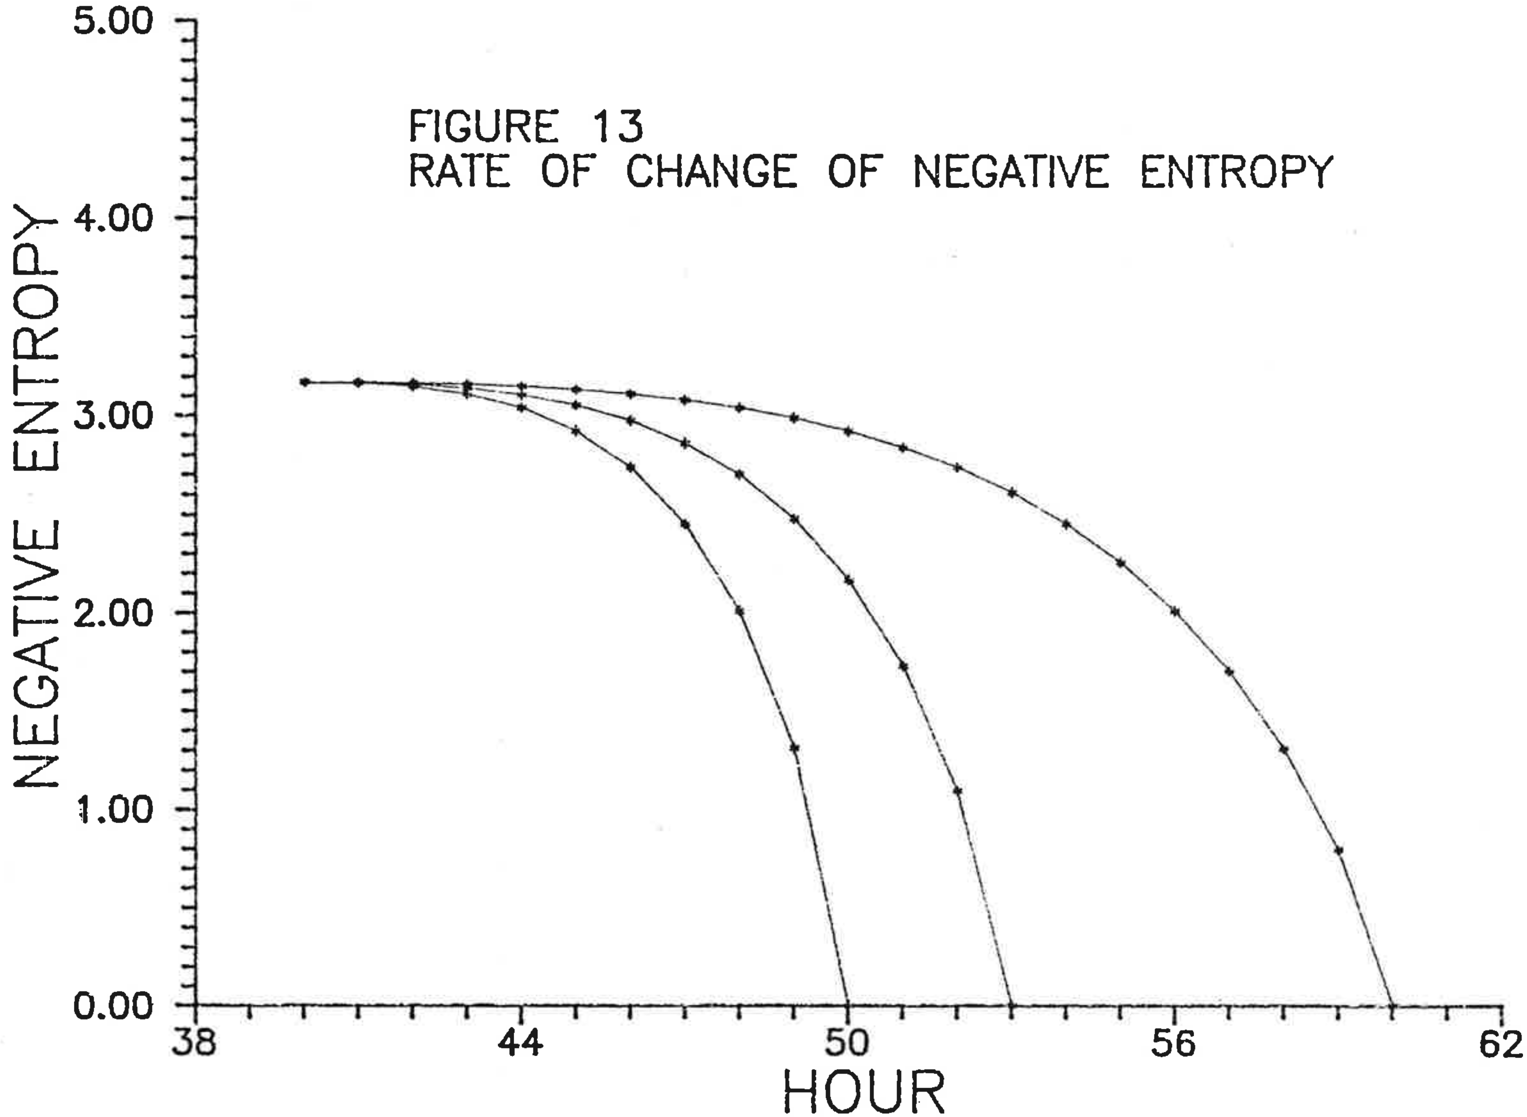

FIGURE 13
RATE OF CHANGE OF NEGATIVE ENTROPY

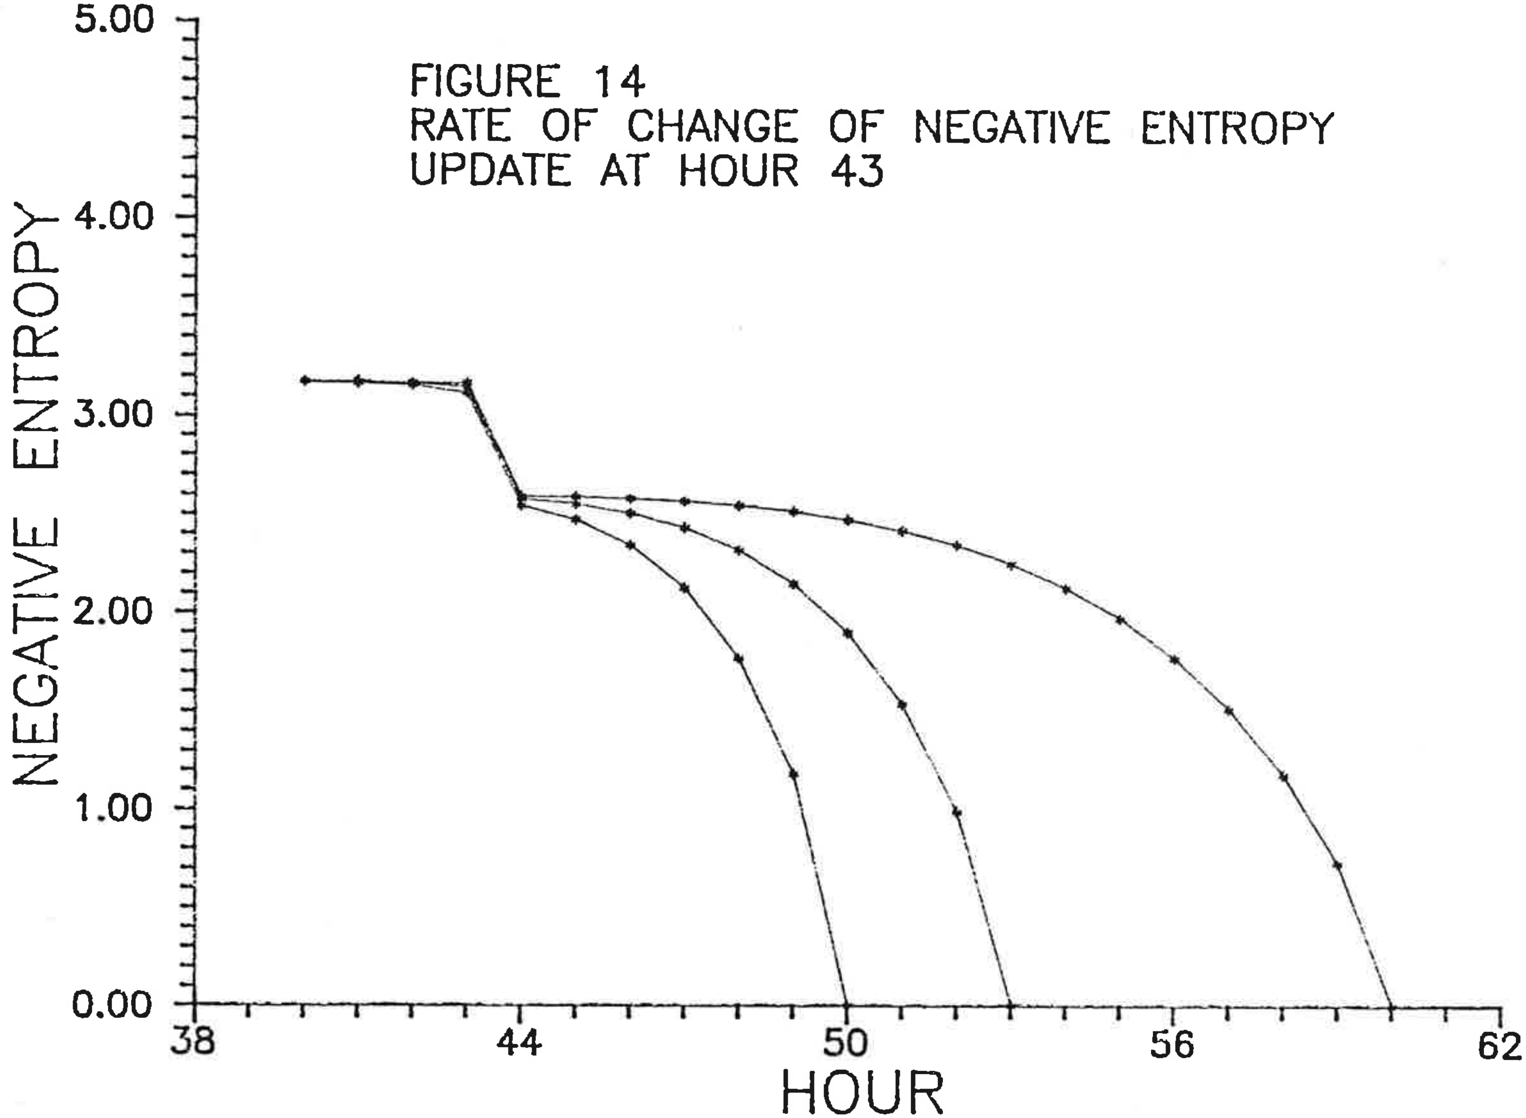
FIGURE 14
RATE OF CHANGE OF NEGATIVE ENTROPY
UPDATE AT HOUR 43
NEGATIVE ENTROPY
5.00
4.00
3.00
2.00
1.00
0.00
38
44
50
56
62
HOUR

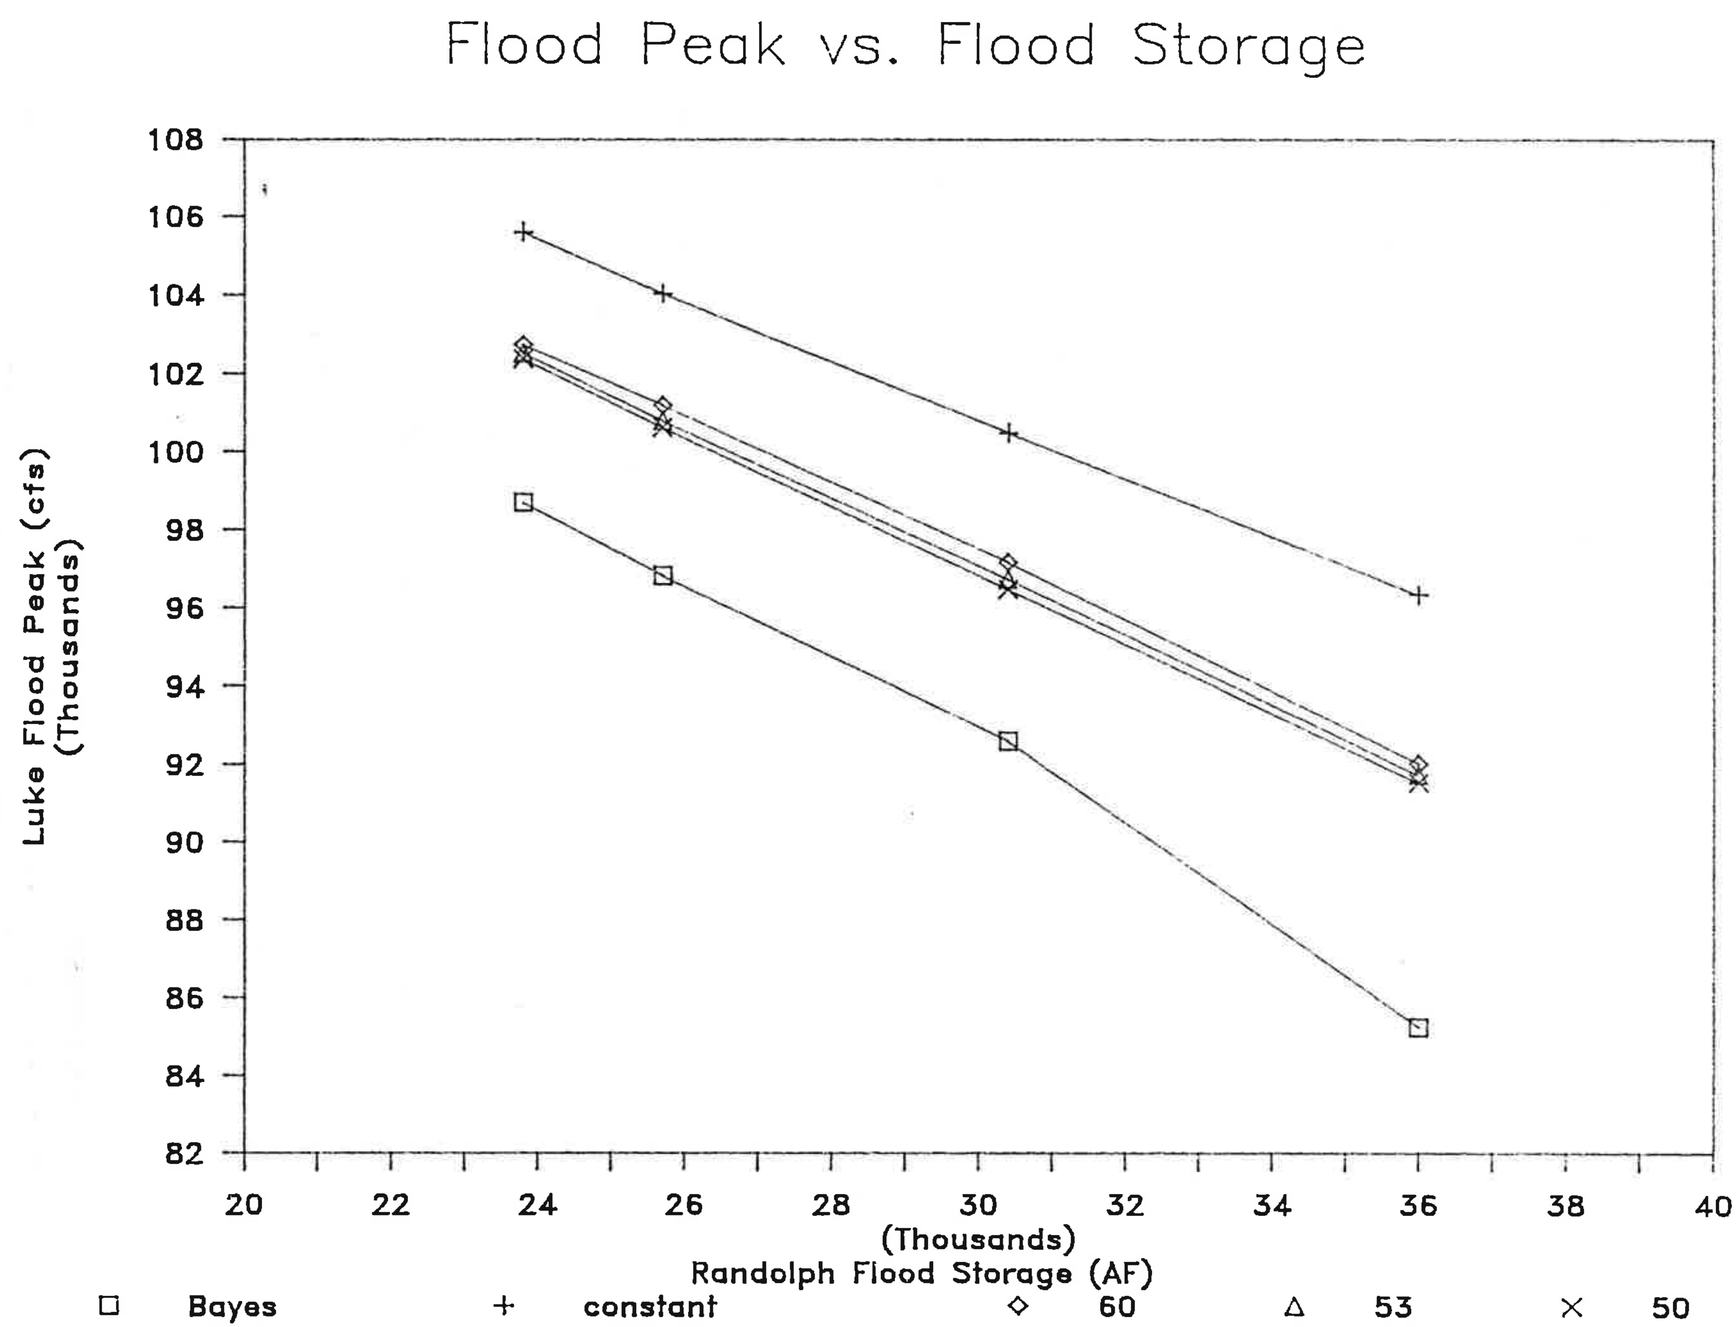

Flood Peak vs. Flood Storage
Luke Flood Peak (cfs)
(Thousands)
108
106
104
102
100
98
96
94
92
90
88
86
84
82
20
22
24
26
28
30
32
34
36
38
40
(Thousands)
Randolph Flood Storage (AF)
Bayes
constant
60
53
50

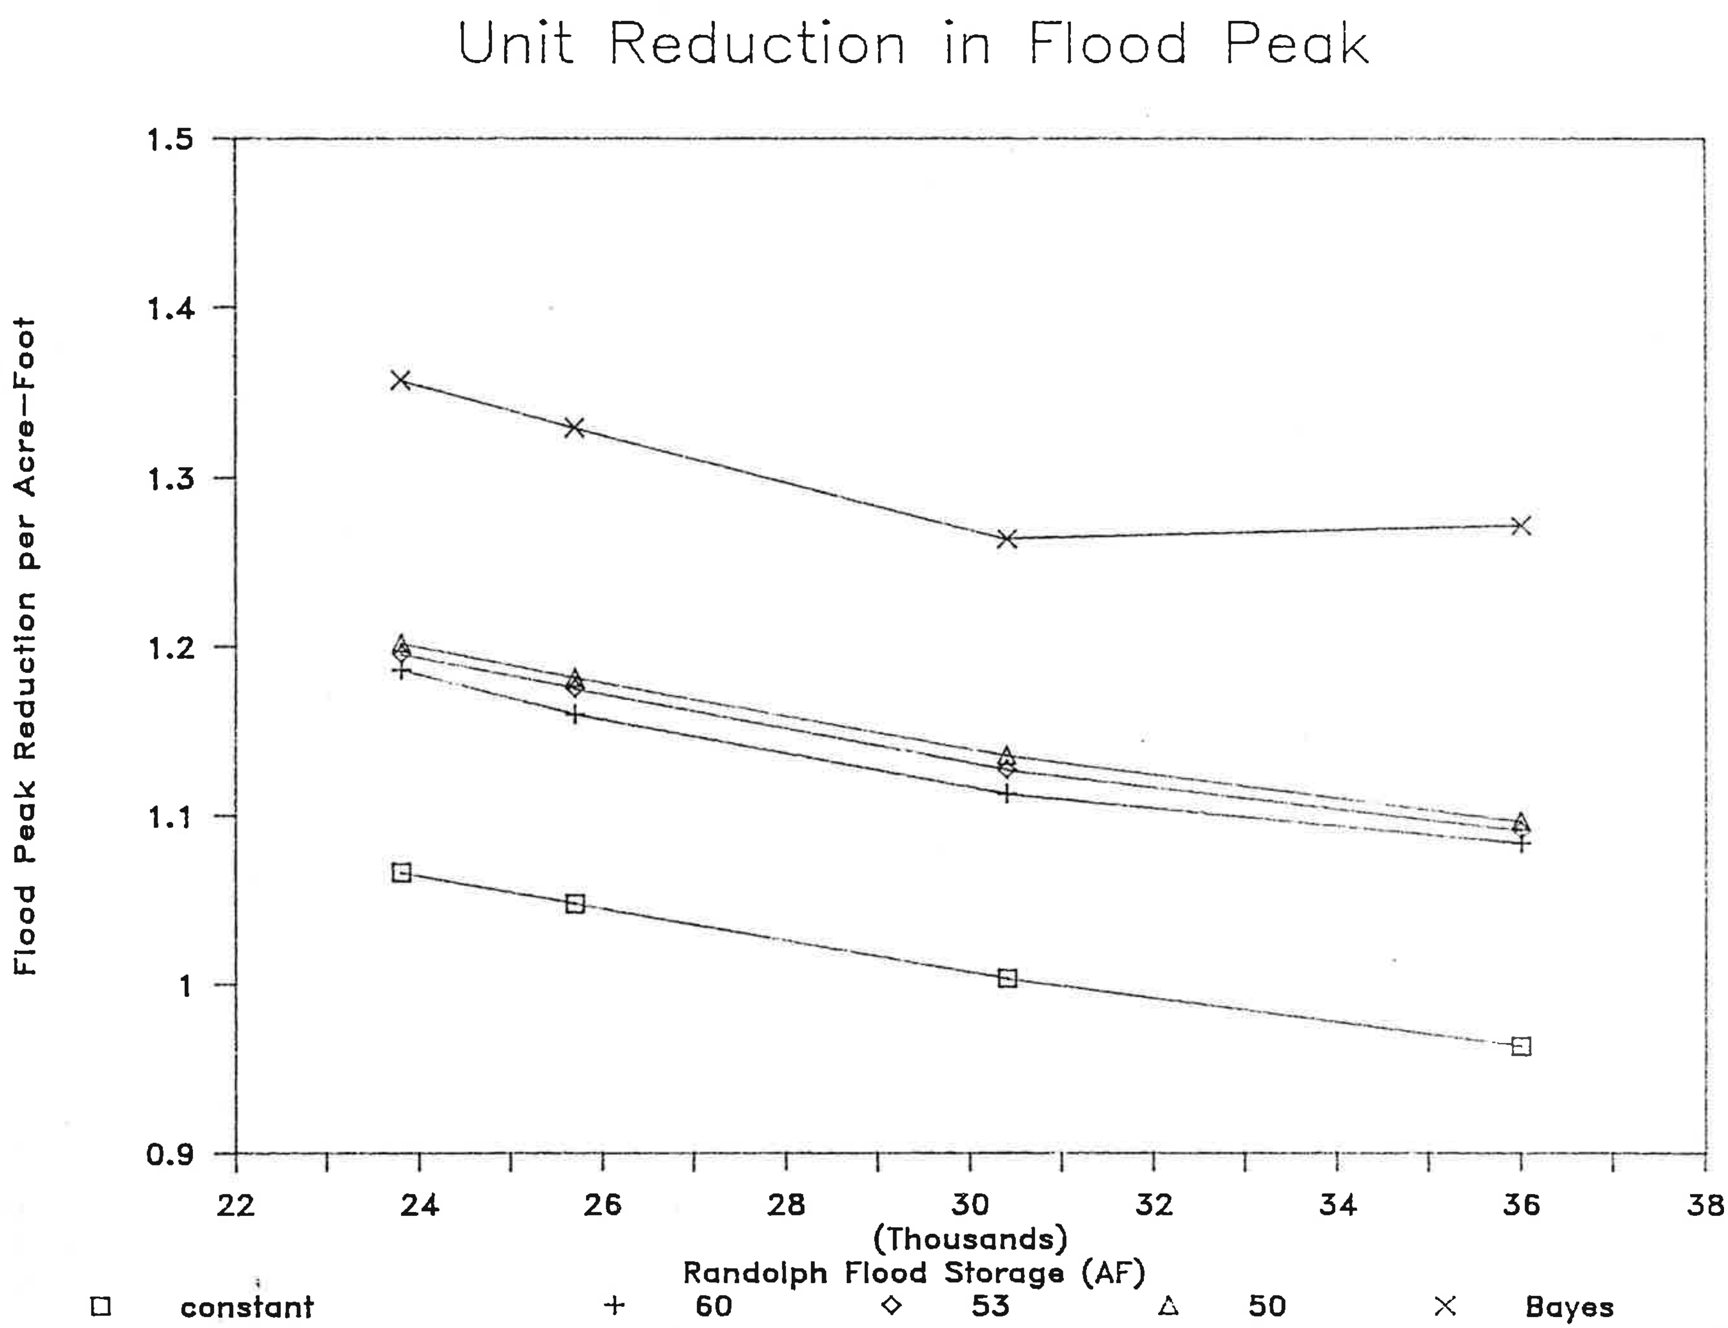
Unit Reduction in Flood Peak
Flood Peak Reduction per Acre-Foot
1.5
1.4
1.3
1.2
1.1
1
0.9
22
24
26
28
30
32
34
36
38
(Thousands)
Randolph Flood Storage (AF)
constant
60
53
50
Bayes

Table 1

Incremental Change in Flood Peak From a
Fixed Reallocation of Flood Storage
Constant Hedging

Flood Storage (AF)	36000	30400	25700	23800
Incremental Change in Flood Storage (AF)	-	5600	4700	1900
Luke Flood Peak (cfs)	96334	101494	104063	105618
Incremental Increase in Luke Peak (cfs)	-	5160	2569	1555

Table 2

Incremental Change in Flood Peak From a Fixed Reallocation of Flood Storage Improving Forecast Skill

Flood Storage (AF)	36000	30400	25700	23800
Luke Flood Peak Constant Hedging (cfs)	96334	101494	104063	105618
Luke Flood Peak Zero Entropy at Hour 60 (cfs)	92001	97162	100549	102760
Luke Flood Peak Zero Entropy at Hour 53 (cfs)	91735	96736	100796	102542
Luke Flood Peak Zero Entropy at Hour 50 (cfs)	91530	96480	100627	102393

<u>Table 3</u>

Incremental Change in Flood Peak From a
Fixed Reallocation of Flood Storage
Improving Forecast Skill &
Truncated Distribution of Runoff

Flood Storage (AF)	36000	30400	25700	23800
Luke Flood Peak Constant Hedging (cfs)	96334	101494	104063	105618
Luke Flood Peak Zero Entropy at Hour 60 (cfs)	92001	97162	100549	102760
Luke Flood Peak Zero Entropy at Hour 60 and Updated Disribution (cfs)	85244	92596	96833	98690

Table 4

Incremental Change in Flood Peak From a Dynamic Allocation of Flood Control Storage Improving Forecast Skill & Truncated Distribution of Runoff

Initial Volume of Empty Flood Storage (AF)	36000	30400	25700	23800
Joint Allocation (cfs)	0	5600	10300	12200
Luke Flood Peak Constant Hedging (cfs)	96334	101494	104063	105618
Luke Flood Peak Zero Entropy at Hour 60 (cfs)	92001	97162	100549	102760
Luke Flood Peak Maximum Forecast Skill (cfs)	85244	92596	96833	98690
Luke Flood Peak Dynamic Allocation Maximum Forecast Skill (cfs)	85244	92005	87630	90050

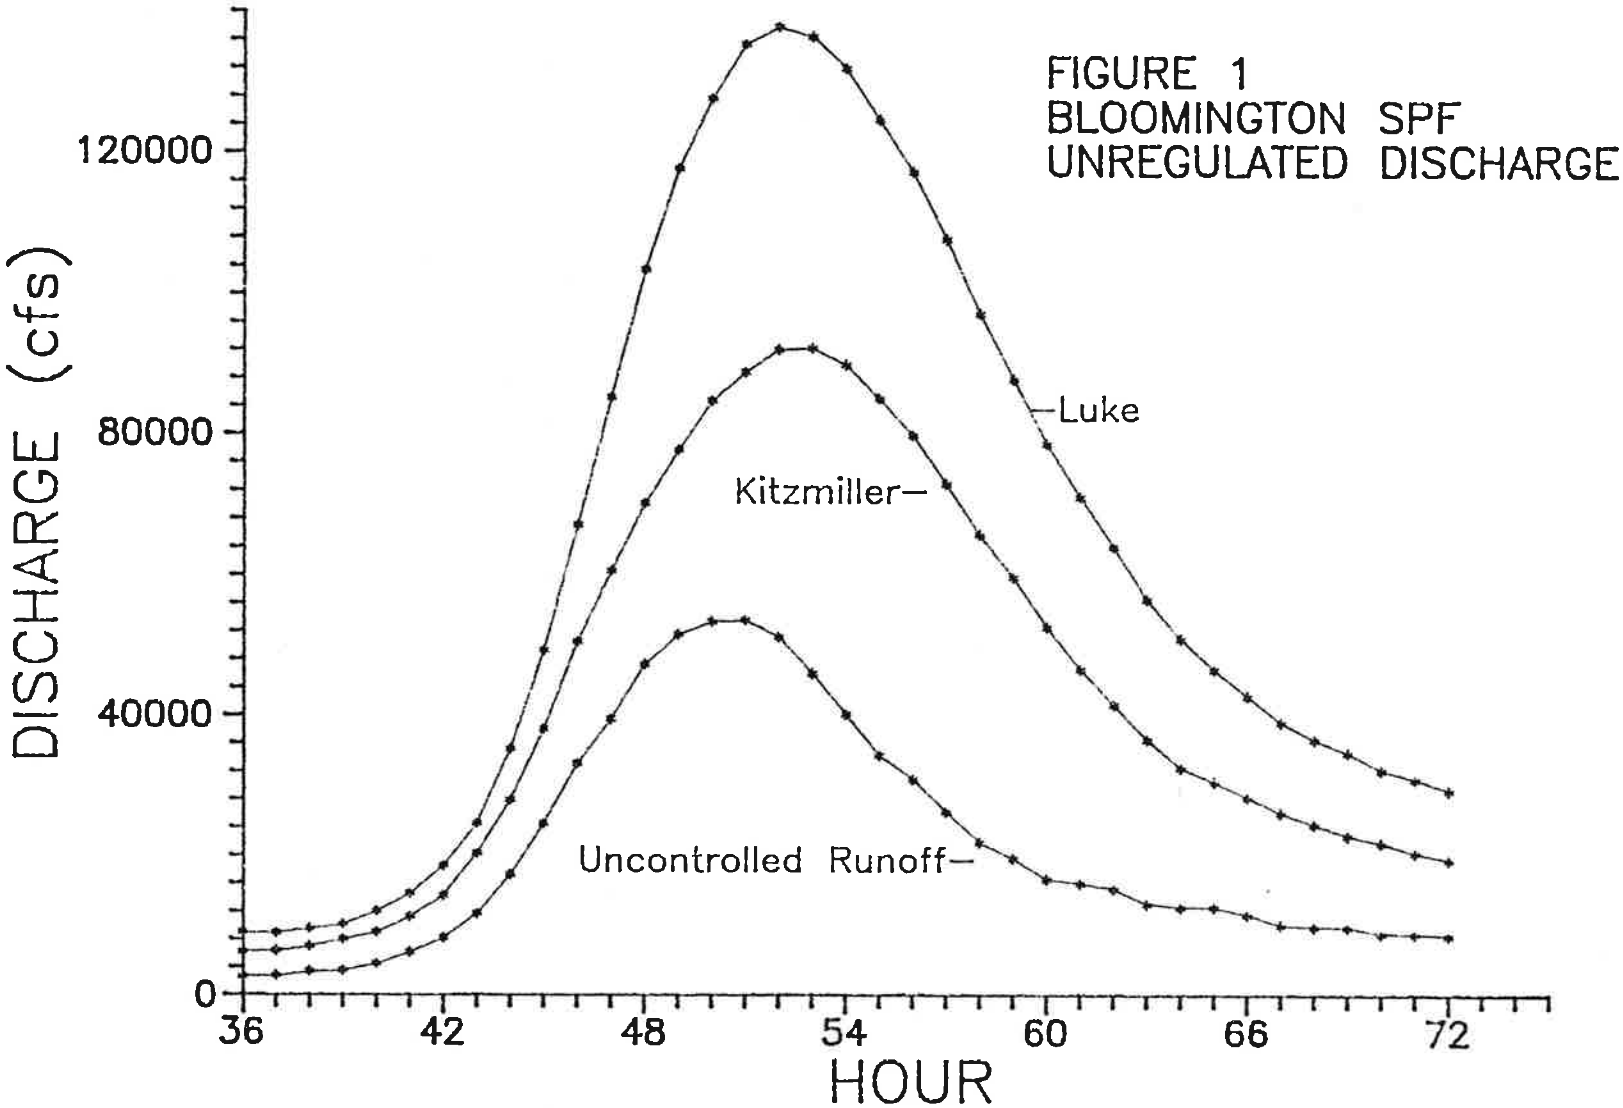
FIGURE 1
BLOOMINGTON SPF
UNREGULATED DISCHARGE
DISCHARGE (cfs)
120000
80000
40000
0
36
42
48
54
60
66
72
HOUR
Luke
Kitzmiller
Uncontrolled Runoff