June 29, 2018

# Aspects of cash balance forecasting.

Cash is essential to government operations; when there is no cash, the government must eventually cease to function. Efficient financial management of the state's resources requires the use of forecasts and projections of cash flows and available assets. Such forecasts provide a lead time for actions to be taken to maintain a smooth trajectory for budget execution and to avoid interruptions in operations arising from cash shortfalls. Cash forecasting is the process of collecting, collating, or creating estimates of cash flows and consolidating these information items into a summary of the government's cash position in advance of any accounting or reports of actual values.

The forecast helps the financial manager gauge the treasury's current position and likely future position over the range of days needed to take action that maintain the treasury's desired balance. In this way, when the forecast is coupled with decision rules on balance adjustment, it plays a role of a servomechanism in the maintenance of adequate cash resources. That is to say, the system detects an out-of-bounds condition and reacts to correct it.

## **The forecasting problem**.

Although the focus of any cash balance forecast must be the sum of funds available to the treasury, the problem will not admit to an easy solution directly. There are many challenges to simply forecasting that total. This is because the daily balance is the net of several cash flows applied to the previous day's closing balance. As such, direct statistical estimation of the total yields an estimate that may have an unacceptable variance. Further, as each of the constituent cash flows that affect the final daily balance have their own characteristics and probability distribution; it is unlikely that these separate random variables can be combined into a probability distribution that conforms to the assumptions needed for many statistical techniques. At a minimum, one may expect heteroscedasticity to be a problem affecting direct estimation.

It is likely more efficient, therefore, to forecast the constituent cash flows and combine them for a net flow. When the net flow is applied to the balance at the beginning of the period of the flow, the closing balance for that period is obtained. It is likely that individual cash flows will have simpler characteristics because they represent a single problem rather than an aggregate of flows. The situation described may be analogous to the following:

An aggregate flow may be too complex to adequately model without error. Shown below is a curve such as might characterize the net cash flow of a government. The horizontal and vertical axes are arbitrary, but they could represent days and currency units respectively. The key point about this curve is that it does not display any regular form. The frequency and amplitude of whatever cycle underlies it appears to shift. To fit this curve to a model may be fairly difficult.

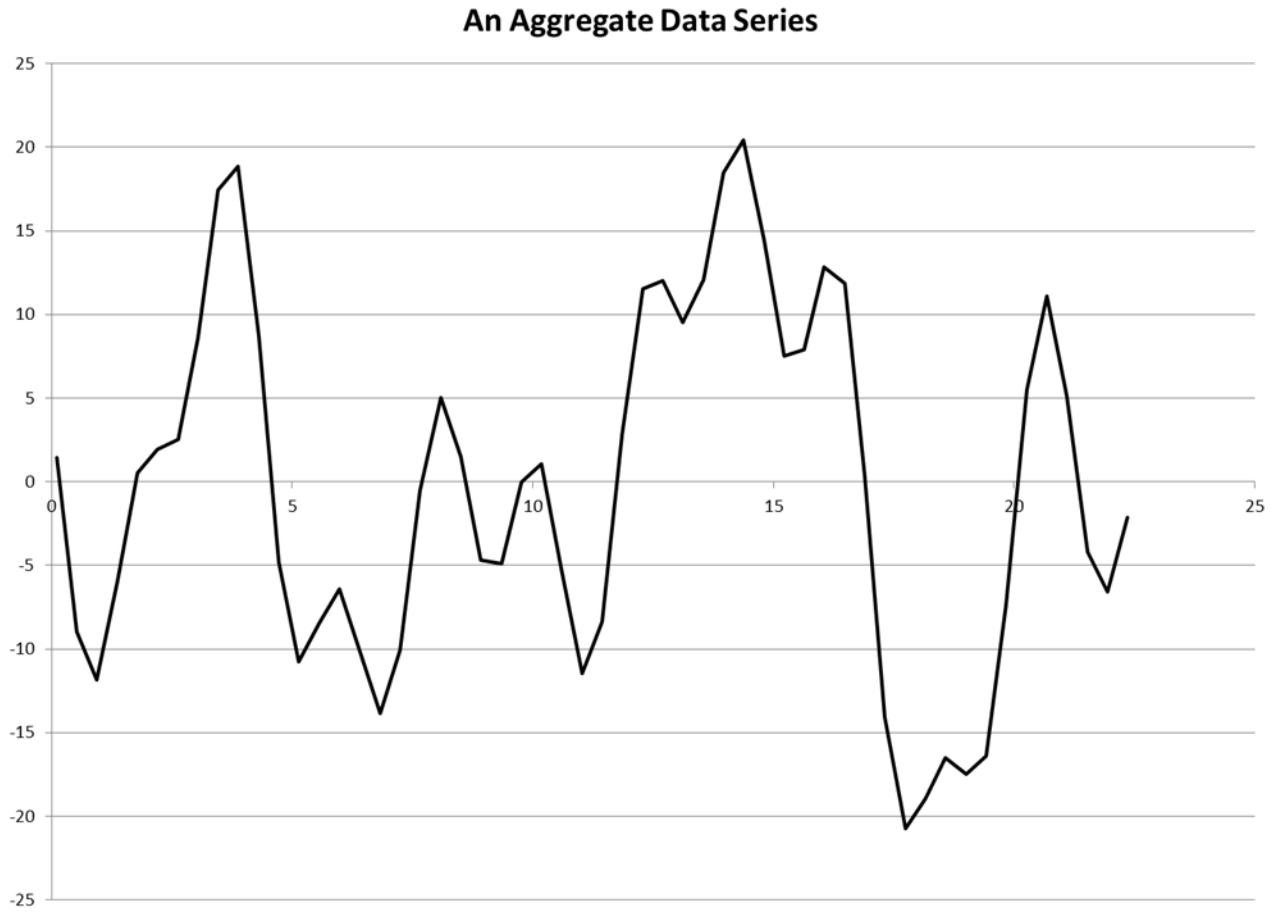


Consider, instead, the problem of forecasting a series of simple functions. Given below are nine separate functions. Each is a sine wave. What varies them is that each has been randomly altered in frequency, amplitude, and phase. Nonetheless, each would pose a much simpler forecasting problem than the aggregate curve above. Their characteristics are obvious to the eye; each is highly predictable.

But this is the key point: the aggregate curve is simply the vertical sum of each of the nine simple functions at each point. In other words, the aggregate is the result of each of the nine simple curves. This leads to the conclusion that a reasonable strategy in forecasting a complex cash flow would be to break it down into constituent cash flows. This should not be surprising because it reflects Fourier's Theorem: a *periodic* function f(x) which is reasonably continuous may be expressed as the sum of a series of sine or cosine terms (called the Fourier series), each of which has specific amplitude and phase coefficients known as Fourier coefficients.

There is no guarantee that the aggregate cash flow will be periodic, although the nature of government operations makes it likely. Nor is it guaranteed that actual individual cash flows will

be as simple as the sine waves given here, but there is a better chance that they will have some elements that are easier to forecast than the aggregate.

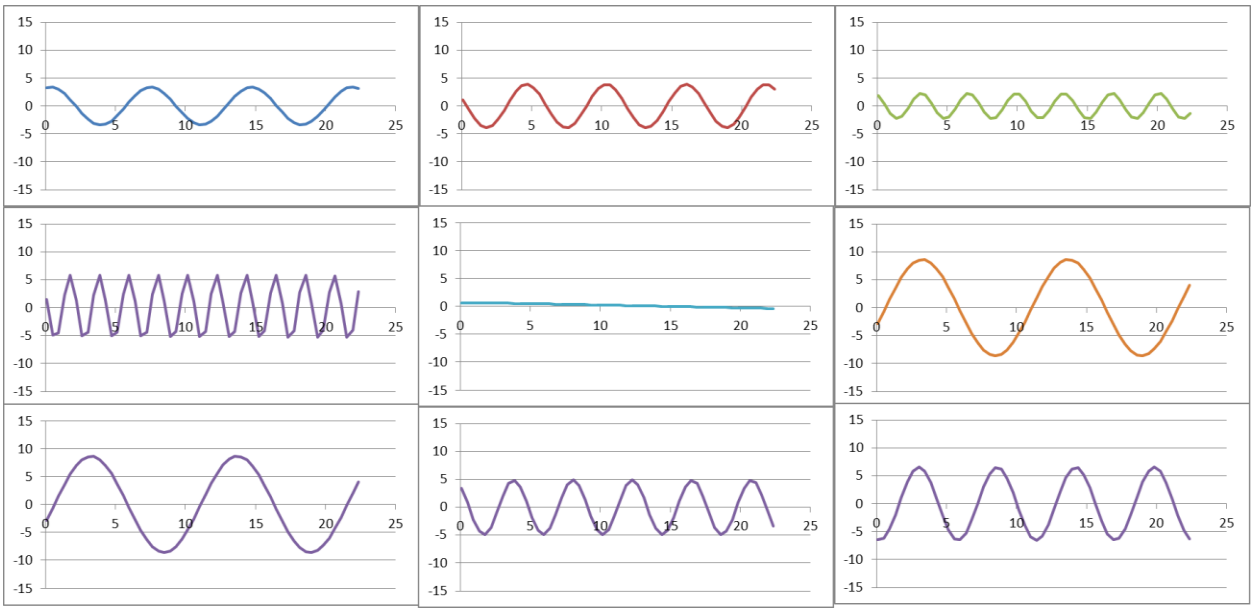

## Critical characteristics of a cash balance forecast

The discussion of how to approach the forecasting problem, by disaggregating cash flows into tractable units however, is not sufficient to define the characteristics of a forecasting model or system. The uses to which it will serve will determine the forecast horizon, the granularity of the estimates, the refresh rate for the estimates, the frequency of issuance, and the coverage of the cash flows.

The *forecast horizon* is measured by the shortest period between the first date of the current forecast and the last date in the forecast. For many countries, both developed and emerging economies, a single instance of the cash forecast is derived as part of the budget approval process by the parliament. This forecast is used to evaluate the realism of the budget proposal; it covers only the fiscal year. If the endpoint remains fixed at the last day of the fiscal year, the forecast horizon actually shrinks to a single day as the fiscal year evolves. This is a problem in emerging economies where the forecast serves no real further purpose than illustrating the budget assumptions. Some European countries derive a forecast with a two-year horizon at its longest. The end point is left unchanged for three quarters before an additional year is added. The horizon effectively varies between 15 and 26 months. The United States Treasury creates a cash balance forecast that extends for nine months at first publication. Although the forecast is issued each day, the same end point is kept until the entire forecast is refreshed three months later. The forecast horizon thus ranges from six to nine months.

The *granularity* of the forecast refers to the smallest unit of measured time in the forecast. For many emerging economies, the forecast only calculates the net flow over each month within the horizon. For a one-year budget forecast, that means the results will list twelve estimates of net

cash flow, one for each month in the fiscal year. More advanced governments may even increase the granularity to weekly totals. Neither of these is adequate for an advanced economy or a government focused on true cash management. All advanced economies generate forecasts of daily granularity.

Financial transactions take place within a day. One day is the basis for calculating interest and float costs. If the essential aspect of cash management is timing, then using a granularity of periods longer than a day masks critical timing issues. The cash management implications, for example, for a government with one large disbursement and a slightly larger receipt each occurring within the month depend on the relative positions of those two cash flows within the month. If the receipt occurs first, all is well. If, however, the disbursement occurs first, the government must find some way to bridge the time period between the outflow and the receipt. Absent any bridge, the expenditure may not be possible. A forecast with granularity of a month will not see the problem before it actually occurs.

The *refresh rate* for the estimates specifies how frequently the entire forecast is revised in its underlying estimated cash flows. The triggering event may be the arrival of fiscal outturn data against which past estimates may be compared. New outturn data offer a means of identifying problems with past estimates. The two should be compared and the pattern of forecast errors inspected for statistical indications of bias or nonrandom patterns. (A distinction should be made between a true refresh, in which the entire forecast in all its estimates and method is reexamined and an update in which a few key values are updated with new data or estimates in the near term are adjusted on the basis of recent events.) Those governments which only produce one run of the fiscal year cash forecast have an annual refresh rate. In the United States, a completely new forecast is created every three months for a quarterly refresh rate.

The *frequency of issuance* refers to how often a new trajectory of cash balances is released to the parties with whom it is shared. The forecast may be released each time there is a change in the trajectory. This may be as often as daily with each addition of actual closing balances or whenever the analysts of the forecasting team make small adjustments to near term estimates based on recent errors. Most developed countries use a daily frequency of issuance because they have transactions systems that can provide daily updates of actual cash flows.

The *coverage* of the estimated cash flows refers to the degree of completeness of the forecast. It may not be practical to include in the estimated flows every conceivable cash flow. Some are too small or too rare to add value to the forecast. It is important to balance the amount of work involved in making the forecast complete against the information content added by the marginal cash flows. The forecast should not be seen as an accounting identity. For comparison, the U. S. cash forecast is based on estimates of fewer than ten receipt flows and about 180 disbursement flows. This is much smaller than the actual number of cash flows that might be identified in the U.S. government.

In developing estimates of individual cash flows, the forecasting unit should look at both the spending and revenue plans submitted by the administrative units or agencies responsible for each flow, but also look at the historic patterns. The inertia of government bureaucracies often means that programs are carried through with a predictable regularity. The transaction

management system of the treasury should provide data series needed to develop a pattern of how individual cash flows actually went in recent years.

**The methodology of disaggregation**

As argued above, the cash balance forecast can be derived by *disaggregating* the significant cash flows that affect the government's cash position. The material cash flows, both disbursements and receipts, are identified. Each is then estimated at the level of granularity of the forecast overall for the period of the forecast horizon. This creates a series of estimated amounts for each of the material cash flows.

For each period in the forecast horizon, the cash flow estimates can be summed to create the expected net cash flow for that period. Expenditures will decrease the net flow; receipts will increase it.

Starting with the most current actual cash balance, the first net flow can be added to it to generate an estimate of the closing cash balance for that period. If the net flow is negative, the expected cash balance will decline. If the net flow is positive, the balance will increase. To this estimated closing balance can be added the next period's net flow to generate a forecast of the balance two periods forward. Add to this forecast the next period's estimated flow to obtain the estimated closing balance. By continuing recursively, adding the next net balance to the estimated previous closing balance, one obtains a trajectory of cash balances through the forecast horizon. This is the cash balance forecast.

The forecast can be updated each period by replacing the first estimated balance with the last actual closing cash balance. The difference between these two values summarizes the total net error in that periods estimated net flow. In replacing the earliest estimated cash balance with the actual balance for that same period, the curve that defines the trajectory of cash balances forward will move rigidly up or down. This approach generates a new trajectory each time the forecast is updated with actual outturn data. The forecast, then, includes the latest information available.

Revision or the forecast should include comparison to actual outturns. Because forecasting is subject to error, it is incumbent upon the forecaster to provide an estimate of the range of that error. By comparing *a priori* estimates with *a posteriori* actuals, the forecasting team can identify faulty bases for forecast information. More importantly, it can get some sense of the types of errors to which its forecast is subject. Information about past errors may also be used to attenuate forecasts and improve future accuracy.