HONG KONG MONETARY AUTHORITY

Research Memorandum 22/2005
November 2005

ADJUSTING FOR THE CHINESE NEW YEAR: AN OPERATIONAL APPROACH

Key points:

- *In macro surveillance work, we often find it difficult to assess the state of the Mainland economy based on unadjusted monthly or quarterly data as the underlying economic trends are obscured by seasonal variations and the effect of moving holidays. Standard procedures exist for seasonal adjustment, but they do not deal with unusual data movements caused by the Chinese New Year (CNY) – a moving holiday in the Gregorian calendar.*

- *In this study we refine seasonal adjustment by taking into account the CNY effect. Two methods are used to pre-adjust a series before the actual seasonal adjustment step: a) taking the average of January and February, and b) using CNY dummies.*

- *The empirical analysis shows that taking into account the CNY effects noticeably improves the quality of seasonal adjustment.*

- *There is, however, no clear winner between the two techniques for adjusting for the CNY effect. Given the relative operational ease of averaging January and February data, we intend to adopt this method in our regular data analysis.*

Prepared by : Chang Shu and Andrew Tsang
External Department and Research Department
Hong Kong Monetary Authority

I. INTRODUCTION

In our regular monitoring of the Mainland economy, we usually analyse a number of major monthly and quarterly macroeconomic indicators in assessing the current economic situation. However, it is invariably difficult to discern economic trends from directly inspecting unadjusted time series, especially for monthly data, as a number of factors obscure the underlying movement. The changes in unadjusted series may reflect seasonal variations which arise due to natural factors (*e.g.* weather patterns), administrative measures (*e.g.* school calendar), social/cultural traditions (*e.g.* fixed holidays such as the National day), and the length of the month or quarter. A further complication is that there are moving holidays which can occur in different calendar months or quarters. One typical example in the Mainland's case is the Chinese New Year (CNY). This is a fixed date in the lunar calendar, but moves around in January and February in the Gregorian calendar which is commonly used throughout the world.

Chart 1 uses some unadjusted Mainland data to illustrate the difficulty in analysing unadjusted monthly series due to the existence of seasonal variations and moving holidays. Panel A shows that exports and industrial production typically fall at the beginning of the year when there are a number of major holidays such as the (Gregorian) New Year and the CNY. Activities then gradually accelerate throughout the year, and tend to become especially buoyant towards the end of December. Retail sales and M0 – which reflects the consumption side of the economy – follow a different pattern. They tend to peak at the beginning of the year around the CNY, and decline sharply after the holiday. They start to rise around the middle of the year with the pace of expansion quickening towards the end of the year. More unadjusted series are plotted in Chart 2.

Chart 1: Selected unadjusted series

Industrial production and exports	Retail sales and M0
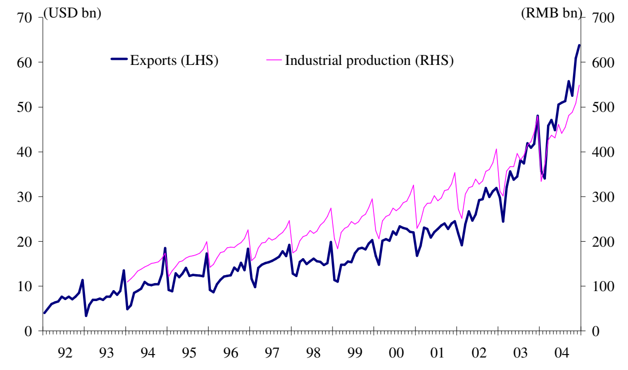	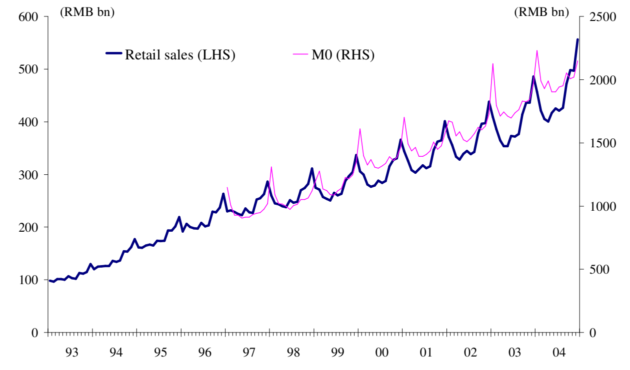

A year-on-year comparison is sometimes used to circumvent problems caused by seasonal variations and moving holidays. However, this practice has several drawbacks. Firstly, the rate of change obtained is based on developments over the preceding 12 months. Recent trends in the data may not be reflected in a year-on-year comparison, making it difficult to identify turning points. Secondly, moving holidays make a year-on-year comparison particularly problematic. For example, different dates of the CNY often lead to large swings of year-on-year growth rates at the beginning of the year. The volatility in economic data makes it very difficult to assess the current state of the economy, making adjustment for seasonal variations and the moving holiday effect an imperative.

Procedures exist for seasonal adjustment, but typically do not remove moving holiday effects. In commonly used packages for seasonal adjustment such as X-12-ARIMA developed by the United States Bureau of Census, it is possible to adjust for holiday effects, but the holiday dates are only defined for the United States. As a result, even though the impact of moving holidays has been noted for other economies and some *ad hoc* methods have been used to make adjustment, few studies have formally addressed the problem. In the Mainland's case, the National Bureau of Statistics releases year-on-year changes of unadjusted series.[1] Even though aware of the problems in using these figures, most major agencies monitoring the economy tend to make assessments based on the year-on-year rates.

[1] The NBS does release seasonally adjusted monetary aggregate data, but with a considerable lag.

This study attempts to refine seasonal adjustment for some Mainland data by taking into account the CNY effect. The aim is to enable more meaningful month-on-month comparisons, as well as to improve the reliability of the year-on-year comparison. Nine major monthly macroeconomic indicators are studied with detailed analysis of industrial production and M0. In the remainder of the paper, Section 2 briefly introduces the techniques that will be used to make adjustment for seasonal and CNY effects. Section 3 discusses in detail the CNY adjustment for industrial production and M0. Section 4 uses one example to illustrate how this refinement in seasonal adjustment helps our understanding of data and assessment of macroeconomic conditions.

II. METHODS OF ADJUSTING FOR SEASONAL AND CNY EFFECTS

The package we use for seasonal adjustment is X-12-ARIMA developed by the Census Bureau in the United States. It has been continually improved since the 1960s, and is used by many statistics agencies and central banks.

The X-12-ARIMA procedure makes adjustment for monthly or quarterly series. It consists of three steps that build upon one another. In the first stage, a regARIMA model is built for the time series. This technique combines the tools of regression analysis with the ARIMA approach to pre-adjust various effects such as outliers, trading day and holiday effects. As the X-12-ARIMA adjustment is best done with a two-sided filter, the fitted time series model is used to simulate data outside the sample period so that two-sided filters can be applied. The second stage of X-12-ARIMA carries out the actual seasonal adjustment which decomposes the pre-adjusted series, *i.e.* the output from the regARIMA step, into three elements – trend, seasonal, and irregular components. The final part of the programme tests the quality of seasonal adjustment.

As noted earlier, although X-12-ARIMA is powerful in seasonal adjustment and takes care of holiday effects for the US calendar, further adjustment is needed for the CNY in the Mainland data to allow meaningful comparisons of monthly data.

There are two possible ways of adjusting for the CNY effect, both of which pre-adjust a series before carrying out the actual seasonal adjustment. One is to simply average the January and February data points. This method is operationally simple, and can be effective. However, its shortcoming is that it removes some dynamics as there might well be variations in economic activity in the two months. The second method involves generating dummy variables for the moving holiday period, and utilises the regARIMA to construct a model which includes these dummies. More details of this approach is given in the Technical Appendix.

There are many ways to evaluate how good a seasonal adjustment is. The seasonally adjusted series should achieve a reduction in volatility. It should not contain regular patterns. It is also desirable if it can be projected forward with a great deal of accuracy.

It should be noted that there is no uniquely best seasonally adjusted series. The decomposition depends on the quality and amount of consistent data available. A good seasonal adjustment can be done only if a time series has an identifiable seasonal pattern. This consists of peaks and troughs which occur in the original series at approximately the same time every year. The more irregular the pattern, the harder it is to separate components for further extracting and removing the seasonal component from a time series. Changes in the series, such as changes in the data collection methodology, cause difficulties in quantifying the components.

III. CNY ADJUSTMENT FOR SOME MAINLAND DATA

This section illustrates the seasonal and CNY adjustment for two monthly data series – industrial production (IP) and M0. The US Bureau of Census' X-12-ARIMA is used for seasonal adjustment, and the CNY dummies are generated by an add-on of the package – *genhol*.[2]

Industrial production

Monthly data for industrial production start from January, 1994. The series spanning 1994.1 and 2004.12 is used for seasonal adjustment.

The first panel in Chart 3 plots the seasonally adjusted industrial production series which has not taken into account the CNY effect. (This series is denoted as 'X12' in charts and tables.) The series is transformed into logarithm in the adjustment process. In the first stage of fitting a regARIMA model, ARIMA(2,1,1)(1 1 0) is chosen based on the *AICC*, and the multiplicative mode is chosen for seasonal adjustment.[3] The trading day effect is not significant possibly because many factories operate on a continuous basis with the exception of a few major holidays. Based on the model, six outliers are identified, all of which are observations in January and February. The series is then backcast and forecast a year using the established model.

Compared to the unadjusted series, IP_X12 – the adjusted series using default X-12 – is smoother. This is reflected by its smaller standard deviation, as well as a smaller average and standard deviation of its monthly changes compared to the original series. However, the chart of IP_X12 shows that even after seasonal adjustment, some regular patterns appear to still exist. In some cases, unusually low values are immediately followed by unusually high values, while in other cases, the opposite happens. This is indicative of the CNY effect, which is more pronounced in the first panel of Chart 4 graphing year-on-year growth rates of the X12 series. As highlighted by the circles in the chart, there are often sharp swings in the growth rates at the beginning of a year. This has motivated work on refining seasonal adjustment by taking into account the CNY effect.

As explained in the previous section, two methods are used for the CNY adjustment. The first method is to take the average of the first two months of the year. The series adjusted using this method is denoted by 'AVE'.

The second method includes in the regARIMA three dummies for before, during and after the CNY (the resulting series is denoted by 'CNY'). In generating the three dummies, the holiday period itself is fixed to be 7 days. The intervals for before and after the CNY are assumed to be the same, and 7, 10, 15 and 30 days are used as possible lengths for the intervals. Based on the *AICC*, a 7-day period is chosen as the optimal interval for the industrial production series for before and after the CNY.

---

[2] These two packages can be retrieved at: http://www.census.gov/srd/www/x12a/.

[3] ARIMA(2 1 1)(1 1 0) means that the model contains one first order difference of the series with a seasonal difference at lag 12. The model can be written as:

$$(1-\theta_1 B - \theta_2 B^2)(1-\theta_3 B^{12})(1-B)(1-B^{12})y_t = (1-\psi B)\varepsilon_t .$$

Chart 3 shows that CNY and AVE are visibly smoother than X12. Sharp swings in the year-on-year growth rates, which are evident in X12, are largely removed from CNY and AVE, especially in the case of AVE (Chart 4).

Tables 2 and 3 demonstrate the greater smoothness of CNY and AVE using different criteria. Table 2 shows that both the final (seasonally adjusted) and trend series from these two methods have smaller standard deviations than those from X12. Table 3 reports that the monthly changes of seasonally adjusted and trend series from CNY and AVE are smaller, and less volatile than those of X12. Comparing the two methods of adjusting for the CNY, AVE appears to achieve a greater degree of smoothness.

Table 2 also shows that the irregular component of X12 is more volatile than that of CNY and AVE2. This is because the holiday factor, *i.e.* the CNY effect, is included in this component. With CNY and AVE correctly attributing the CNY effect to the holiday factor, the variations in the irregular component become smaller. This is also demonstrated visually in Charts 5-7.

One useful feature of a seasonally adjusted series is that it contains useful information for projecting forward. We run a horse race to test the predictive power of the three seasonally adjusted series. We fit an ARIMA model for individual series using data up to 2003.12, and then project forward for the next 12 months. In comparing the forecasts and series adjusted based on the full sample, CNY and AVE are shown to produce more accurate projections than X12 with the CNY series being the most predictable series (Table 4).

Overall, the two methods of taking into account the CNY clearly improved seasonal adjustment. Among the two, taking the first two months' average may be a better way to adjust for the CNY effect for the industrial production series. It achieves the greatest reduction in volatility, and is effective in removing nearly all regular patterns in the original series. Its forecasting performance is not as good as using the CNY dummies, but shows significant improvement compared with X12.

M0

Monthly data for M0 date back to 1997.1, and the time period between 1997.1 and 2004.12 is used for seasonal adjustment.

Unlike industrial production, in all the regARIMAs for the three adjustment procedures (X12, CNY and AVE), trading day effects are significant for M0. Demand for cash tends to fall on Mondays and Tuesdays, but rises on Sundays. The optimal interval for before and after the CNY is chosen to be 15 days based on the *AICC*.

The first panels of Charts 8 and 9 show that the seasonal adjustment by the routine X-12 is highly unsatisfactory. Noticeable patterns exist in the adjusted series, and year-on-year growth rates show sharp swings at the beginning of the years in five out of the seven-year data span.

The adjusted series of CNY and AVE and their year-on-year growth rates are graphed in the second and third panels of Charts 8 and 9. Both these charts and Tables 2-3 show that these two methods have improved the seasonal adjustment to a certain extent. CNY and AVE are smoother than X12, and there are fewer large swings in their year-on-year growth rates compared with X12. The average monthly changes and their standard deviations of CNY and AVE and their trends are smaller than their counterparts derived from X12. CNY and AVE also have a better forecasting performance than X12.

Nevertheless, CNY and AVE series still appear to be volatile, especially AVE. Although they have helped remove a number of irregular movements in the original series, they also seem to introduce some unusual movements, *e.g.* a blip in 2003 in CNY. Comparing the two series, AVE is more volatile than CNY, and has retained some regular patterns, albeit less pronounced than X12.

Overall, although CNY and AVE have improved performance of seasonal adjustment of M0 in terms of removing some regular seasonal patterns, volatility reduction and forecasting ability, the results are not entirely satisfactory. This is possibly due to a short sample period. There are only seven years of data available for M0, when often a span of 10 years is needed in order to establish a seasonal pattern.

Similar analysis has been carried out for a number of other major monthly macroeconomic series. Table 4 compares the performance of three adjustment procedures.

IV. HOW DOES THE WORK HELP WITH MACROECONOMIC SURVEILLANCE?

The introductory section discussed the difficulties the moving holiday bring to data analysis. This section uses one example to demonstrate how the refinement of taking into account the CNY effect helps with the assessment of the macroeconomic situation.

Industrial production is a series that is particularly affected by the CNY. When the CNY does not fall in the same month as in the previous year, the growth differential (measured in year-on-year rates) between January and February can be as wide as seventeen percentage points at its biggest. One typical example is that the CNY was in February in 2003, but in January in 2004. The unadjusted IP series shows that, after a strong growth in the region of 17% in the second half of 2003, the year-on-year growth of industrial production suddenly decelerated to 7.2% in January 2004, but then surged to 23.2% in February. The sequential growth momentum calculated based on the seasonally adjusted series without taking into account the CNY also shows a sharp downward blip in January 2004. The annualised 3-month-on-3-month rate fell by more than 14 percentage points compared with the previous month, or 13 percentage points lower than in the fourth quarter of 2003. The turn of 2004 was an important time in macroeconomic adjustment of the Mainland economies. The authorities took some tightening measures in 2003 aiming at cooling down the economy. Relying solely on the above figures alone could have led an erroneous assessment that the economy had slowed down sufficiently, or even too sharply, and thus further policy actions would not be needed.

However, taking into account the CNY changes the picture. In carrying out the adjustment, we use data up to the end of 2003, which were data available at the time to policymakers and China observers. The seasonal factors obtained for 2003 are then applied to adjust 2004 figures. The resulting year-on-year growth rate was 14.6% for January, 2004 -- much higher than the rate that does not account for the CNY effect. The annualised 3-month-on-3-

month growth rates also point to a robust growth momentum, as industrial production expanded by 17.5%, which matches the strong growth in the previous period. This set of figures provides a more accurate picture of macroeconomic conditions at the time. They would signal to the authorities that the growth momentum had not abated at the beginning of 2004, and further policy measures were required to cool down the economy.

V. CONCLUDING REMARKS

This paper has illustrated how the Chinese New Year – a moving holiday in the Gregorian calendar – has complicated data analysis. We believe that the CNY effect should be routinely checked before carrying out seasonal adjustment for Chinese data, and if it exists, adjustment needs to be made accordingly.

We use two methods to take into account the CNY effect, both of which pre-adjust a series before the actual seasonal adjustment step. They are: a) taking the average of January and February, and b) including CNY dummies in the regARIMA. We compare in detail the performance of the three procedures – the routine X12, and the two methods accounting for the CNY – in adjusting industrial production and M0, but also report the adjustment for a number of other major monthly macroeconomic series.

The comparison shows that taking into account the CNY effects clearly improves the quality of seasonal adjustment in that they reduce volatility and produce more accurate projections. This refinement in seasonal adjustment is shown to aid us in making more accurate assessments of macroeconomic conditions in the Mainland.

However, there is no clear winner between the two techniques of adjusting for the CNY. Given the relative operational ease of averaging January and February data, we intend to adopt this method for our regular data analysis.

Technical Appendix: Adjusting for Moving Holidays

This appendix explains how to use holiday regressors to take into account holiday effects in a seasonal adjustment procedure. This method involves generating dummy variables for the moving holiday period, and utilises the regARIMA to construct a model which includes the dummies. Different methods have been used for creating such dummy variables. One focuses on the CNY period itself, and uses a single regressor. This is proposed by Bell and Hillmer (1983) to model the effect of the Easter holiday. Assume the holiday affects the economy for a total length of τ days, and the effect is the same for each day during this interval. Let τ_t denote the number of days in month t that belong to this interval. The holiday regressor, $H(\tau,t)$, is then defined as:

$$H(\tau,t) = \frac{\tau_t}{\tau}.$$

For example, assume a 10-day period for the CNY. The CNY dummy is set to be zero for all the months except for January and February. Then the days within the two months are counted. If, for example, one day of the 10-day period falls into January while the other 9 into February, the CNY dummy takes the value of 0.1 and 0.9 for the two months respectively.

While one single regressor may be sufficient for modelling Easter in western economies, several might be needed for modelling other holidays. As we observed in the previous section, production slows for the periods before and after the CNY, and completely stops during the holiday, while consumer behaviour may also differ in the three periods. In this case, three regressors – $H_1(\tau,t)$, $H_2(\tau,t)$ and $H_3(\tau,t)$ – can be used for before, during and after the holiday.

However, the value of τ is often unknown. To select the most suitable τ, we use a modified Akaike Information Criterion proposed by Hurvich and Tasai (1989) which is shown to have better performance for small samples. The criterion is defined as:

$$AICC = -2\log likelihood + 2p\frac{1}{1-\dfrac{p+1}{T-12D-d}},$$

where p is the number of estimated parameters, D is the order of seasonal differencing, d order of regular (non-seasonal) differencing. The model with the smallest $AICC$ is preferred.

References

Bell, W.R., and Hiller, S.C. (1983). 'Modelling time series with calendar variations,' *Journal of the American Statistical Association*, **78**, 526-34.

Hurvich, C.M. and Tsay, C.L. (1989). 'Regression and time series modelling in small samples,' *Biometrika*, **76**, 297-307.

Table 1: Dates for the Chinese New Year

January	1982, 1987, 1990, 1993, 1998, 2001, 2004
February	1981, 1983, 1984, 1985, 1986, 1988, 1989, 1991, 1992, 1994, 1996, 1997, 1999, 2000, 2002, 2003, 2005
January-February	1995

Table 2: Standard deviations of component series from the seasonal adjustment procedure

	Seasonally adjusted series	Trend	Irregular component	Seasonal factor	Trading day effect	CNY effect
Industrial production						
X12	0.3495	0.3500	0.0149	0.0873	N.A.	N.A.
CNY	0.3489	0.3478	0.0081	0.0902	N.A.	0.0137
AVE	0.3488	0.3487	0.0050	0.0905	N.A.	N.A.
M0						
X12	0.2457	0.2456	0.0210	0.0770	0.0051	N.A.
CNY	0.2434	0.2431	0.0126	0.0624	0.0057	0.0077
AVE	0.2434	0.2445	0.0106	0.0598	0.0061	N.A.
Retail sales						
X12	0.4166	0.4149	0.0102	0.0623	N.A.	N.A.
CNY	0.4166	0.4160	0.0072	0.0627	N.A.	0.0051
AVE	0.4167	0.4166	0.0059	0.0619	N.A.	N.A.
CPI						
X12	0.5550	0.5547	0.0033	0.0129	N.A.	N.A.
CNY	0.5550	0.5549	0.0035	0.0129	N.A.	0.0005
AVE	0.5550	0.5549	0.0034	0.0130	N.A.	N.A.

Table 3: Volatility of final seasonally adjusted and trend series

	Seasonally adjusted series		Trend	
	S1	S2	S1	S2
Industrial production				
X12	0.9831	2.4614	0.9847	0.3273
CNY	0.9717	1.2564	0.9704	0.3023
AVE	0.9673	0.7806	0.9687	0.2981
M0				
X12	0.8567	3.6130	0.8633	1.2099
CNY	0.8362	1.8311	0.8417	0.3685
AVE	0.8173	1.7941	0.8416	0.9429
Retail sales				
X12	1.1056	1.6309	1.1057	0.6297
CNY	1.0547	1.1949	1.0498	0.6119
AVE	1.0523	0.9886	1.0531	0.6100
CPI				
X12	0.5159	0.8192	0.5166	0.6766
CNY	0.5154	0.8342	0.5157	0.6787
AVE	0.5168	0.8122	0.5168	0.6742

Notes: S1 refers to the average of month-on-month changes, and S2 is the standard deviation of month-on-month changes.

Table 4: Summary of regARIMA modelling

	Data available from	ARIMA model	Forecasting error (%)
Domestic demand/production			
Industrial production	1994.1		
X12		(0 1 2)	7.53
CNY		(0 1 2)	6.48
AVE		(0 1 2)	6.91
Fixed asset investment	1995.1		
X12		(2 1 2)	37.06
CNY		(1 1 0)	31.15
AVE		(2 1 2)	37.06
Retail sales	1993.1		
X12		(1 1 2)	6.65
CNY		(2 1 1)	3.48
AVE		(1 1 2)	3.25
External trade			
Exports	1992.1		
X12		(1 1 2)	6.30
CNY		(1 1 2)	5.68
AVE		(1 1 2)	3.34
Imports	1992.1		
X12		(1 1 2)	14.94
CNY		(0 1 2)	14.75
AVE		(1 1 2)	4.47
Monetary development			
M0	1997.1		
X12		(0 1 2)	5.37
CNY		(0 1 2)	4.47
AVE		(2 1 0)	4.97
M2	1997.1		
X12		(2 1 1)	1.62
CNY		(2 1 1)	1.50
AVE		(2 1 1)	1.84
Credit	1997.1		
X12		(2 1 2)	2.06
CNY		(2 1 2)	1.67
AVE		(1 1 1)	0.85
Price			
Consumer price index	1981.1		
X12		(2 1 2)	0.69
CNY		(2 1 2)	0.92
AVE		(2 1 2)	0.64

Chart 2: Non-seasonally adjusted series

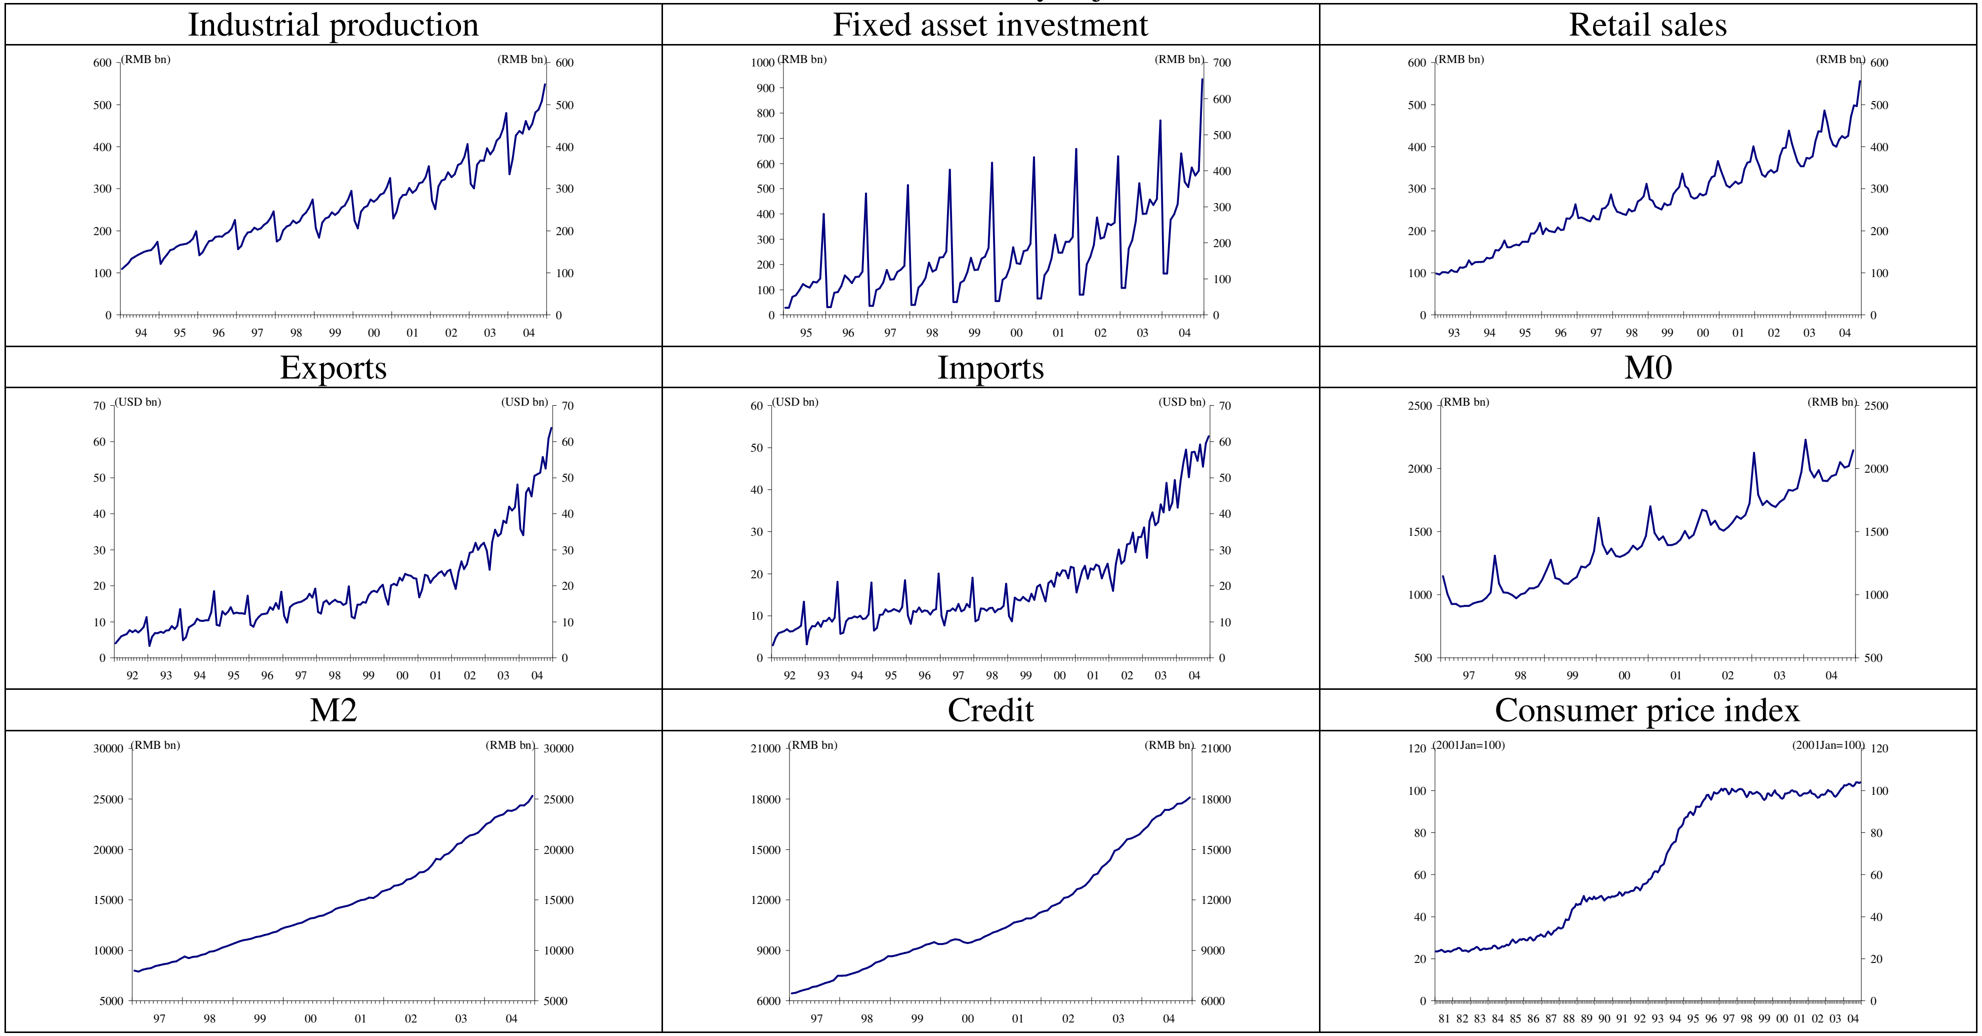

Chart 3: Seasonally adjusted series of industrial production

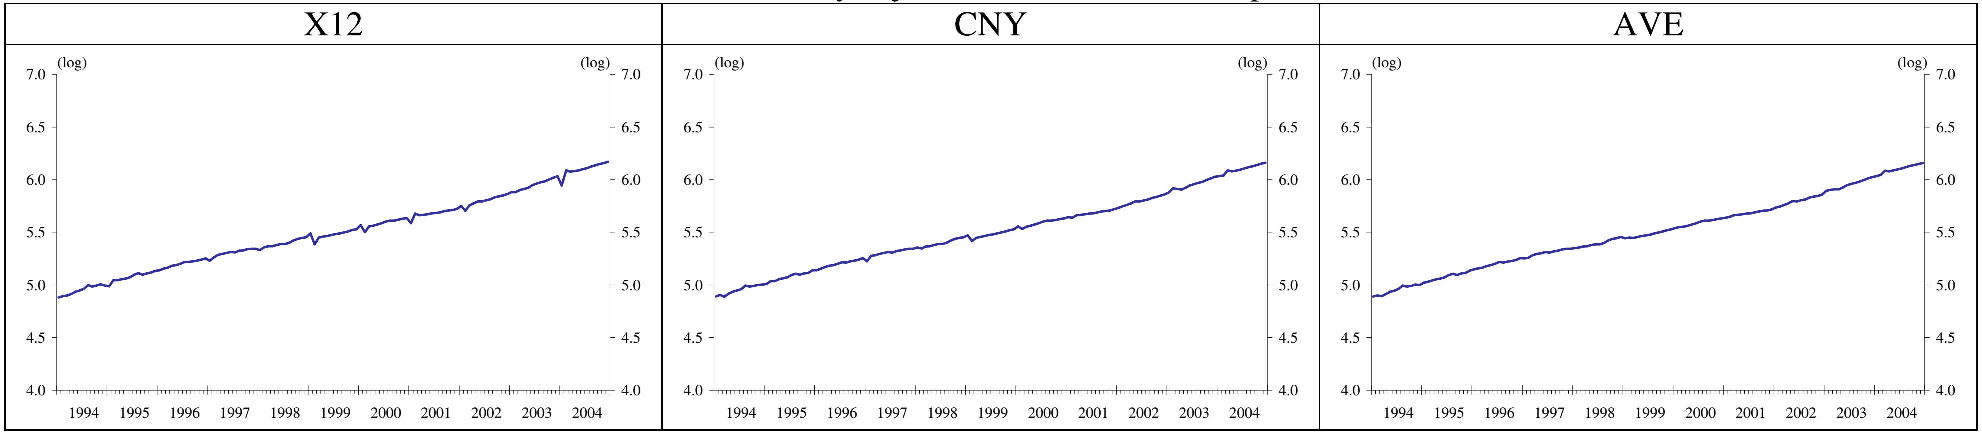

Chart 4: Year-on-year growth of industrial production

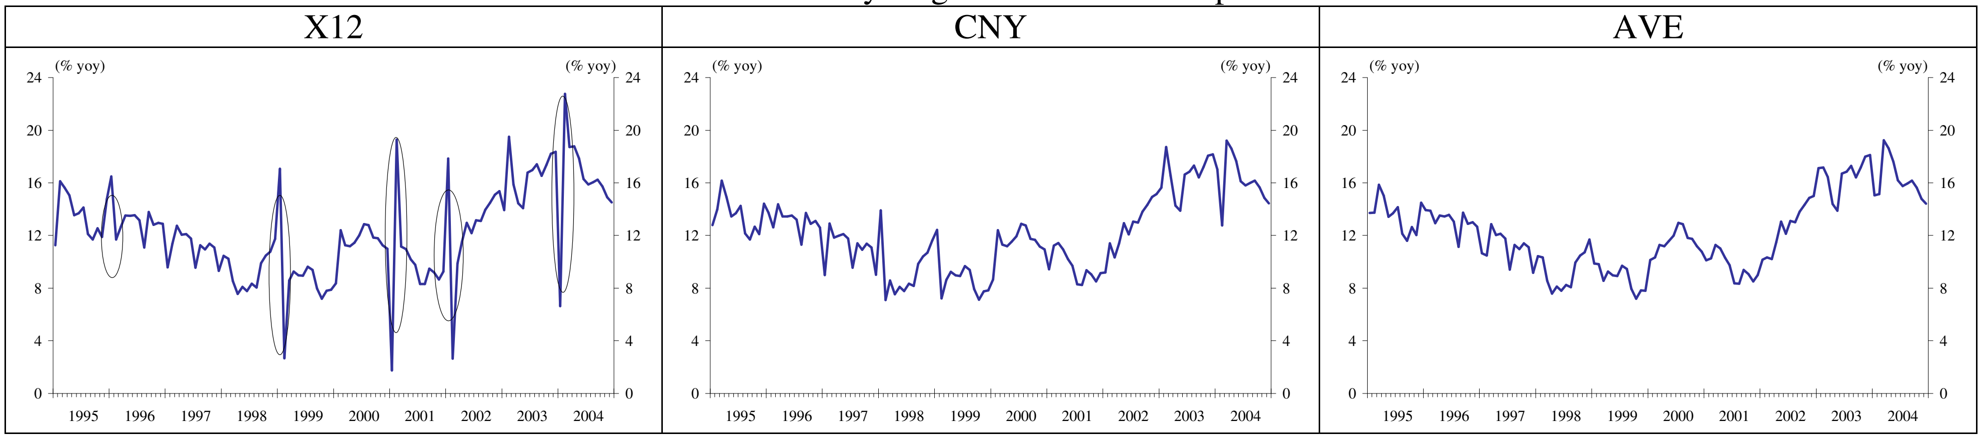

Chart 5: Irregular components of industrial production

X12	CNY	AVE
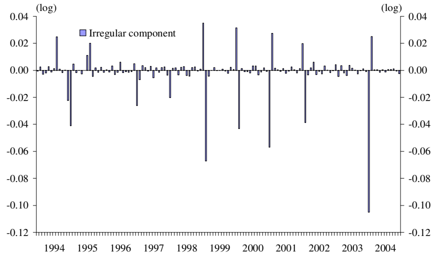	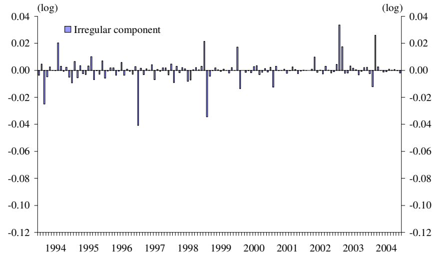	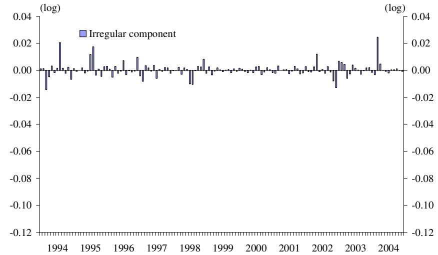

Chart 6: Seasonal factors of industrial production

X12	CNY	AVE
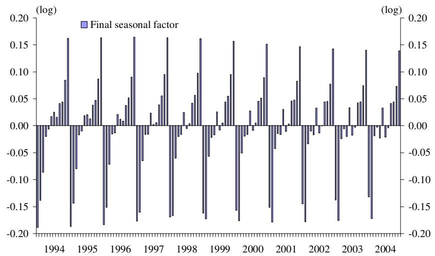	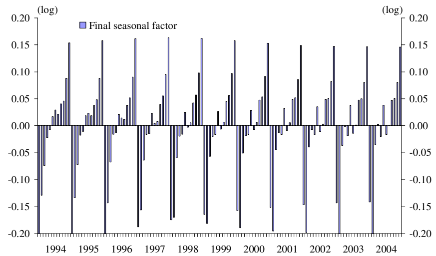	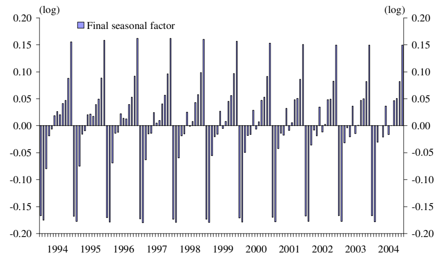

Chart 7: Holiday factors of industrial production

X12	CNY	AVE
N.A.	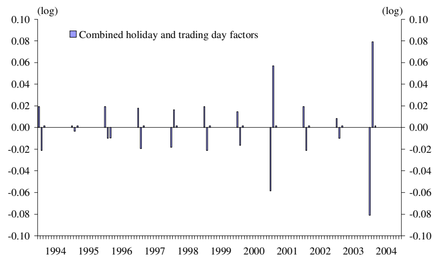	N.A.

Chart 8: Seasonally adjusted series of M0

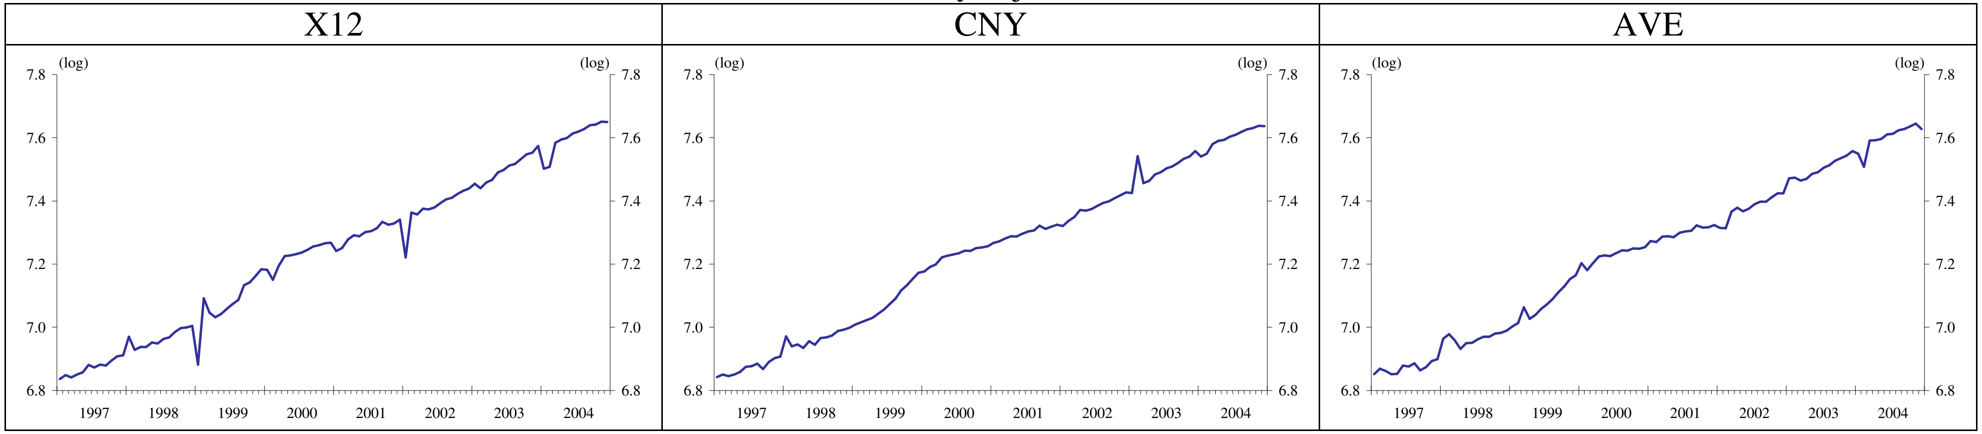

Chart 9: Year-on-year growth of M0

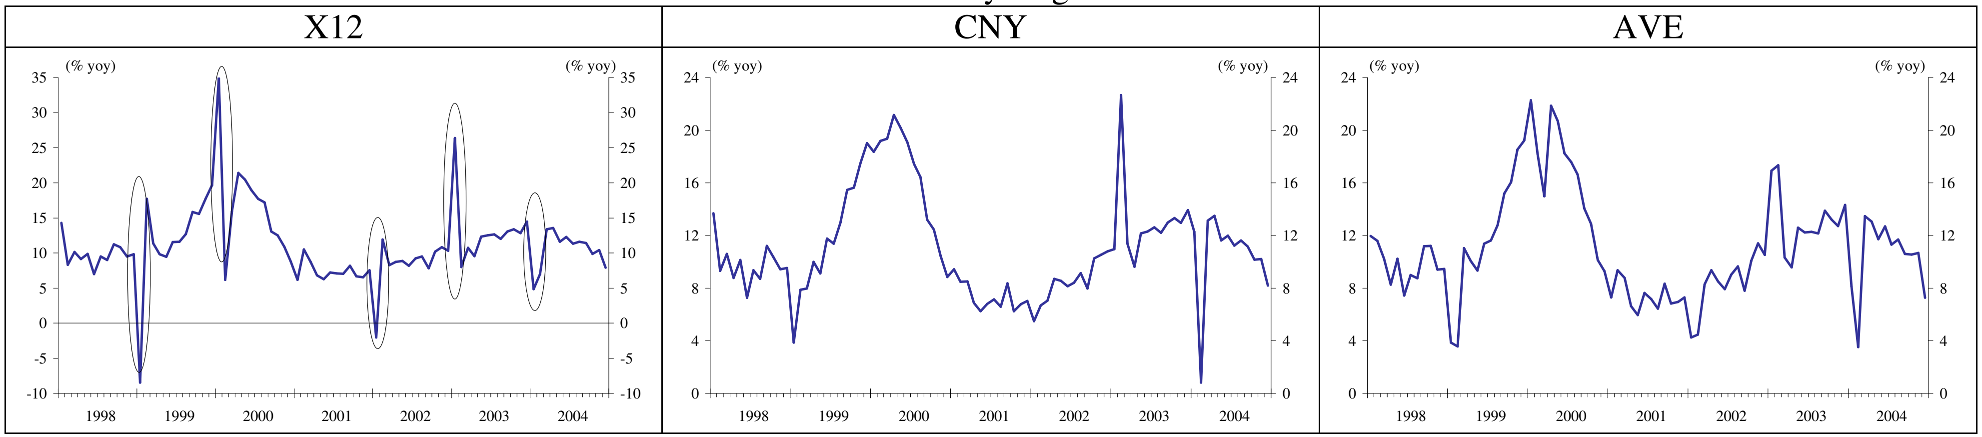

Chart 10: Irregular components of M0

X12	CNY	AVE
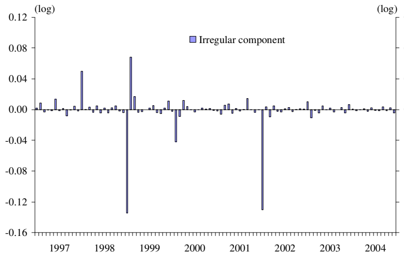	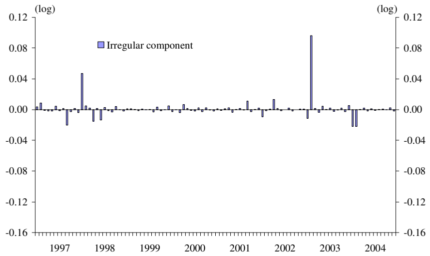	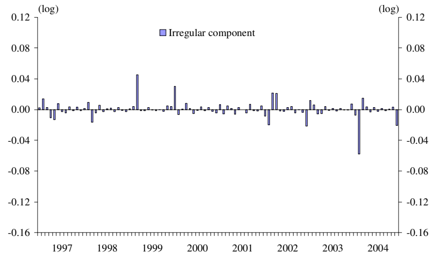

Chart 11: Seasonal factors of M0

X12	CNY	AVE
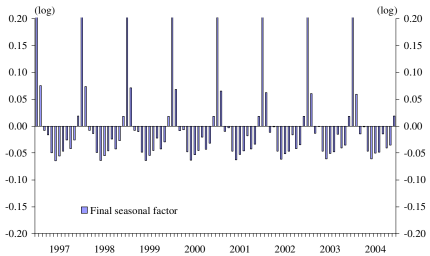	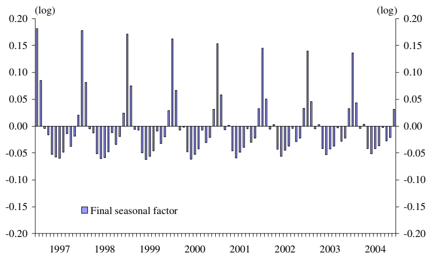	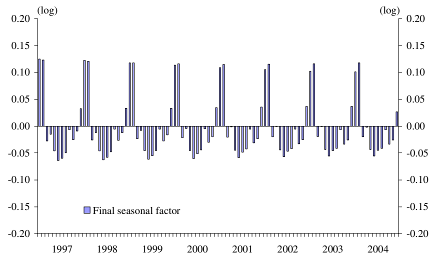

Chart 12: Holiday factors of M0

X12	CNY	AVE