

# Time Series student project: An analysis of historical gold prices

David Saul Davis

June 23, 2012

## 1 Introduction

Throughout human history gold has been a symbol of wealth. Its rare physical properties (e.g., its resistance to oxidization) and appropriate abundance (i.e., rare enough to make it hard to get, but common enough that it can be widely used) have made it perhaps the single most important precious metal in the history of human finances. Historically, many currencies were backed by gold, driving central banks to amass large stockpiles. While this practice has fallen out of paper as national governments all over the world realize the advantages of fiat money. However, there remains a deep liquid market for gold on several exchanges throughout the world. For this project, I will examine quarterly gold prices since 1979, and determine the if there is a simple underlying structure.

## 2 Data

For this project, I obtained quarterly values for the price of gold from the World Gold Council ([http://www.gold.org/investment/statistics/gold\\_price\\_chart](http://www.gold.org/investment/statistics/gold_price_chart)). These prices were listed in nominal United States dollars per ounce. The quarterly data on this site began in 1979, giving 31 years of data. The raw data are shown in the top panel of Figure 1.

Because there has been significant inflation over the last thirty years, it is important to convert nominal dollars to real dollars. To this end, Consumer Price Index (CPI) data was obtained from the Federal Reserve Bank of St. Louis (<http://research.stlouisfed.org/fred2/series/CPIAUCSL/downloaddata>). The year which should be used as the baseline is arbitrary. I therefore transformed all nominal dollars into 2012 dollars. The inflation-adjusted data is shown in the bottom panel of Figure 1. The data is complete (i.e., there were no missing values), and a visual inspection did not reveal any apparent erroneous values. Hence the inflation adjustment was the only correction made to the data.



Figure 1: Both raw (top) and inflation adjusted (bottom) gold prices (averaged over a quarter) since 1979.

### 3 Differences

A quick glance at Figure 1 will reveal that the nominal price of gold does not appear to be stationary. While from the mid 1980s to the early 2000s the price was relatively flat, in the early 1980s and late 2000s the price was extremely volatile. Since 2000 there has been strong steady growth. As an aside, this suggests that it may be fruitful to examine the data in two sections, dividing the data into a low- and high-volatility regimes. Obviously inflation has been a component of the price increase. Examining the inflation-adjusted growth rates reveals that in real dollars there was in fact a price decline for the first two decades of this time series. Many commodities are thought to be reasonably well modeled by a log-normal process. Assuming that this is true for gold, it would be reasonable to test if the first differences of the logarithm of the price appears to be stationary. The 1<sup>st</sup> differences are shown in the top panel of Figure 2. The first differences appear to be stationary. However, there is a slight upward trend in the data. There are many more positive values after the year 2000 than before. This trend is marginal, and clearly could be due to random fluctuations. While we would prefer a model that uses first differences (due to our affinity to parsimony), it seems prudent to explore 2<sup>nd</sup> differences as well. The second differences are shown in the bottom panel of Figure 2. The 2<sup>nd</sup> differences are clearly stationary, so differences of order higher than two will *not* be considered.

### 4 Autocorrelation functions

According to Chan and Cryer, a good place to start exploration of a time series is by examining both the autocorrelation (ACF) and partial autocorrelation (PACF) functions. This is because, if the time series can be well modeled by either a pure moving average (MA) or autoregressive (AR) model, it will have a characteristic ACF and PACF. For a pure AR model, the ACF will decay exponentially for lags beyond its order, with the PACF will have a sharp cut off. For a MA model, the situation is reversed—the PACF will exhibit a decay, while the ACF will show a sharp cutoff. Of course, when using real data, the behavior is often ambiguous. Furthermore, if the model is mixed (i.e., ARMA), this procedure will be ambiguous. Both the ACF and PACF were therefore calculated for both the 1<sup>st</sup> and 2<sup>nd</sup> differences of the gold prices using Cryer and Chan's `acf` and `pacf` routines in R. The results are shown in Figures 3 and 4.

The ACF in Figure 3 shows the first two lags are significant, while the lag at 6 is marginally significant. Lags of higher order show a sine pattern with decaying magnitude. This is characteristic of an AR process. However, the fact that the PACF does not cut off after a lag of two suggests that an AR does not tell the whole story. The PACF shows that similar lags are significant. The fact that neither the ACF nor the PACF show sharp cutoffs suggests that this may be a mixed model. The 2<sup>nd</sup> differences (shown in Figure 4) show similar behaviour. While the difference is not clear, the 2<sup>nd</sup> differences could possibly

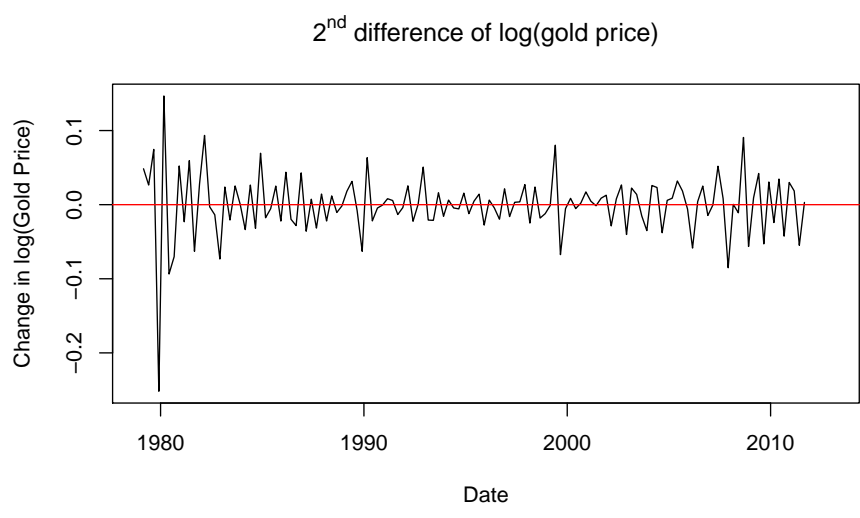
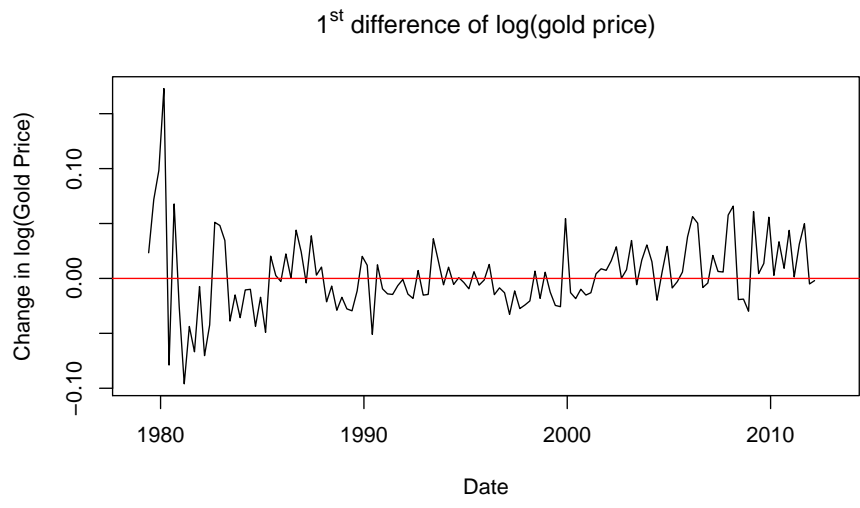


Figure 2: The 1<sup>st</sup> (top) and 2<sup>nd</sup> (bottom) differences of gold prices.

AR/MA	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0	x	x	0	0	0	x	0	0	0	0	0	0	0	0
1	x	0	0	0	0	x	0	0	0	0	0	0	0	0
2	x	x	0	0	0	0	0	0	0	0	0	0	0	0
3	x	x	x	0	0	0	0	0	0	0	0	0	0	0
4	x	0	x	x	0	0	0	0	0	0	0	0	0	0
5	0	x	x	x	0	0	0	0	0	0	0	0	0	0
6	x	0	0	0	0	0	0	0	0	0	0	0	0	0
7	x	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 1: The EACF of the 1<sup>st</sup> difference of gold prices (generated using Cryer and Chan’s `eacf` routine). The vertex of the imperfect triangle suggests the order of the model is ARMA(1,1).

AR/MA	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0	x	0	0	0	0	x	x	0	0	0	0	0	0	0
1	x	x	0	0	0	0	0	0	0	0	0	0	0	0
2	x	x	0	0	0	0	0	0	0	0	0	0	0	0
3	x	x	0	0	0	0	0	0	0	0	0	0	0	0
4	x	x	0	0	0	0	0	0	0	0	0	0	0	0
5	x	x	x	x	0	0	0	0	0	0	0	0	0	0
6	0	x	x	0	0	0	0	0	0	0	0	0	0	0
7	0	0	x	x	0	0	0	0	0	0	0	0	0	0

Table 2: The EACF of the 2<sup>nd</sup> difference of gold prices (generated using Cryer and Chan’s `eacf` routine). The vertex of the imperfect triangle suggests the order of the model is ARMA(0,1).

show a cutoff in the ACF and a tailing off in the PACF. This would suggest a MA model. However, the behavior is not clear, and I therefore will explore mixed ARMA models.

Because the analysis using the ACFs and PACFs was inconclusive, and suggested mixed ARMA models, we will use the extended ACF (EACF) to simultaneously determine the values  $p$  and  $q$  to characterize the model. Cryer and Chan’s R routine `eacf` was used to generate Tables 1 and 2. These tables suggest an ARMA(1,1) and ARMA(0,1) model would be appropriate for the first and second differences respectively.

## 5 Fitting ARIMA models

Having established a rough guess of the possible underlying models, we will begin to test the fits of the suggested models as well as “nearby” models.

1<sup>st</sup> difference of log(gold price)




Figure 3: The ACF and PACF of 1<sup>st</sup> differences of gold prices. The ACF shows the first two lags are significant, while the lag at 6 is marginally significant. The PACF shows similar lags are significant, suggesting an ARMA model.

2<sup>nd</sup> difference of log(gold price)




Figure 4: The ACF and PACF of 2<sup>nd</sup> differences of gold prices. The ACF shows that lag 1 is significant, and lags 6 and 7 are marginally significant. The PACF shows that lags 1, 2, and 6 are significant.

p	d	q	LL	AIC
0	0	0	-948.72	1899.44
0	1	0	-781.98	1563.96
0	1	1	-781.62	1565.25
1	1	0	-781.46	1564.92
1	1	1	-780.00	1564.00
0	1	2	-777.13	1558.26
2	1	0	-777.62	1559.65
1	1	2	-775.44	1556.89
2	1	1	-777.37	1560.73
<b>2</b>	<b>1</b>	<b>2</b>	<b>-773.07</b>	<b>1554.14</b>
0	2	0	-815.80	1631.60
0	2	1	-777.80	1557.60
1	2	0	-788.12	1578.24
1	2	1	-777.59	1559.18
0	2	2	-777.66	1559.31
2	2	0	-781.43	1566.85
1	2	2	-776.77	1559.53
2	2	1	-774.08	1554.15
2	2	2	-773.77	1555.55

Table 3: The log-likelihoods and AIC of various ARIMA( $p, d, q$ ) models. Note the ARIMA(2,1,2) model has the best fit.

$$\begin{array}{cccc}
\phi_1 & \phi_2 & \theta_1 & \theta_2 \\
1.59 \pm 0.04 & -0.95 \pm 0.36 & -1.52 \pm 0.12 & 1.00 \pm 0.15
\end{array}$$

Table 4: Fitted parameters for the ARIMA(2,1,2) model.

The previous section suggested that the underlying model may be either a ARIMA(1,1,1) or ARIMA(0,2,1). Model parameters were estimated using Cryer and Chan's `arima` R routine. The log-likelihoods and Akaike's Information Criterion (AIC) are tabulated for various values of  $p$ ,  $d$ , and  $q$  in Table 3.

Of the models tested here, the best fit (according to the AIC) was the ARIMA(2,1,2) model. This is slightly surprising as this was not indicated by the EACF tests (shown in Table 1). However, these test were inconclusive, so it is not shocking that they did not predict the optimal model. The fitted parameters are shown in Table 4.

In the ACF in Figure 3, there was a significant lag at 6. This will obviously not be addressed by the ARIMA(2,1,2) model. It is difficult to believe that there is a fundamental reason why the price of gold should be linked to that six quarters ago, but not to that three, four, and five quarters ago. I will therefore not attempt to deal with the significance of the peak at lag six, but rather chalk it up to random fluctuations.



## 6 Residuals

Having decided on a model to fit, i.e., ARIMA(2,1,2), and having obtained the best fit parameters (as shown in Table 4), we now must examine the residuals to see if there are any glaring problems. Figure 5 shows a plot of the residuals (top panel), and a quantile-quantile plot in the bottom panel. The residuals are clearly heavy-tailed. This is obvious both from the residual plot and the QQ plot. However, the heavy tails are dominated by measurements from the previously-identified high-volatility regimes in the early 1980s and late 2000s. The residuals are relatively well behaved.

The next step is to analyze the residuals to see if they are consistent with white noise. Hence, the ACF and PACF were calculated for the residuals. This is displayed in Figure 6. The strongly significance of lags one and two have clearly been removed. There is still a marginally significant lag at six. As mentioned before this is almost certainly attributable to random fluctuations.

Finally, as a check, the `arima` routine was used to check “nearby” models (i.e., nearby to the ARIMA(0,0,0) model for white noise). All models checked at AIC larger than that for the ARIMA(0,0,0), and we can therefore conclude that the residuals are randomly distributed, and the model is a good fit.

## 7 Conclusion

Quarterly gold prices since 1979 were analyzed using R, and Cryer and Chan’s TSA package. The original data was clearly not stationary, so differences were analyzed. It was unclear if the first or second differences would provide a better fit, so models for both were considered. As a starting point, the ACFs and PACFs were examined. If the underlying model was either a pure AR or pure MA model, the ACFs and PACFs would have characteristic shapes. However, these tests were inconclusive, suggesting a mixed ARMA model would likely be required. Examining the EACFs, it there were hints that a ARIMA(1,1,1) or ARIMA(0,2,1) model would be the best fit. However, after examining the AIC for many nearby models, the best fitting model was found to be an ARIMA(2,1,2) model. The residuals, though heavy-tailed, were not found to have any interesting structure, and it was concluded that they are consistent with white noise.

### ARIMA(2,1,2) residuals





Figure 5: The residuals of gold prices fit with an ARIMA(2,1,2) model. Note the periods of high volatility, particularly in the early 1980s, are plainly visible in the top plot. The QQ plot shows that the residuals are heavy tailed.

### Analysis of residuals




Figure 6: The ACF and PACF of the residuals of gold prices fit with an ARIMA(2,1,2) model. The lags of 6 are found to be marginally significant. This is likely due to random noise.