


A Tale of the Guilder's Fat Tail

One of the oldest and longest running financial contracts once was the guilder-pound foreign exchange rate. It started trading in the famous Amsterdam coffee houses during the golden age of the Dutch Republic. The contract witnessed the stormy Anglo Dutch relationship, the century long fall from grace of the pound and the makeover of the guilder into the euro. Not only is the history of this contract a fascinating tale, it also gives ample opportunity to demonstrate our current understanding of financial markets and the risk of sharp movements due to the fat tail phenomenon in particular.

1. Historical Development

We are in the fortunate possession of data on the guilder-pound contracts from the very beginning in 1609. The Dutch economist Posthumus meticulously hand collected data from old newspapers (*Amsterdamsche beurscourant*) and we have completed these by more recent data. Since the ECB was basically modeled after the German and Dutch vision on monetary policy, we felt free to extrapolate these data to the present. The four hundred years of history are summarized in the Figure 1 below. The figure has some gaps due to lost issues and periodic absence of trade. During the Napoleonic years, for example, Holland was sealed off from trade with Brittany.



The relationship between the guilder and the pound was relatively stable until the First World War, due to the precious metals basis of the currencies of both countries. Several historic events led to temporary disturbances. In 1717, the House of Commons, succumbing to public pressure, fixed the value of the (golden) *guinea* to 21 *shilling*. This was above the market price of silver. The oldest law in economics, Gresham's law (1519-1579), predicts that bad money drives out

good money. By overpricing gold, the UK in effect adopted a gold standard as the relatively more precious silver coins were hoarded and quickly disappeared from circulation. The guilder's value, though, remained based on the price of silver. The depletion of Brazil's gold mines in the last part of the eighteenth century, led to an increase in the price of gold over time and with it the pound ascended in value. When England suspended the convertibility of the pound in 1798 to quickly raise finances for its wars against Napoleon, the pound fell again in value. The gold standard was restored around 1820. The Dutch currency remained silver-based until 1873. After 1875 the Dutch also adopted the gold standard, which further stabilized the guilder-pound foreign exchange rates.

The pound depreciated over 15% during the course of the First World War, but returned to its original value in the twenties due a combined effort of the FED and the Bank of England. The reintroduction of the gold standard required a massive creation of dollars and a deflationary policy in the UK, policies which ultimately led to the crash in 1929. As the depression wore on, the UK parted from the gold standard in 1931, leaving the Dutch official pound reserves in shambles. The Dutch stubbornly stuck to the Gold standard with hugely disadvantageous terms of trade. For the open Dutch economy this finally proved too much in 1936, when the last countries finally deserted the metallic standard. These events led to massive fluctuations in the pound-guilder exchange rates during the inter-war years.

At the end of the Second World War the Bretton Woods system of fixed but adjustable exchange rates was established. The Dutch guilder was re-valued in 1961 and the pound was devaluated in 1967. Once the Bretton Woods system collapsed and exchange rates started to float freely, the pound nosedived further down south due to the lax monetary policy in the UK. While the Dutch Central Bank was situated at a healthy distance from the government, the Old Lady from Threadneedle Street was a willing servant of the ministry of finance. Even the iron lady could not resist the temptation and at times reverted to easy finance.


The Dutch influence on Britain's monetary history is considerable. King William in 1694 founded the Bank of England to raise finances for the war against France. In an interesting twist of history, the pound was finally rescued from its century long slide by another Dutch economist. In 1997 Willem Buiter advised the incoming labor government to finally cut the ties between the bank of England and the Chancellor of the Exchequer. The Old Lady gained her independence and the pound finally stabilized against the guilder during the closing act of this historic contract. More recently, the old fashioned euro monetary policy derived from the Bundesbank focus on monetary aggregates and constrained expansion of credit proved to be more stable than the inflation target cum parliamentary debate approach of the Anglo-Saxons. After ten years, the pound is almost a par with the euro. The history of this contract can be succinctly summarized by this Mondriaanesque abstraction.



2. Fundamentals

To which extent do the developments of the exchange rate reflect the developments of the relative purchasing power of the two currencies? Consider the relationship between the price ratios of the two countries. Figure 3 shows the ratio of the price indices over the period between 1620 and 2000, with 1900 as the base year. Two things stand out: The relative decline in prices at the end of the Dutch Golden Age, and the relatively high prices in Britain during the past fifty years.

The figure also includes the exchange rates over the same period. Economists speak of purchasing power parity if the prices of one country translated into a foreign country's currency by means of the exchange rate are identical. International trade and migration arbitrages price differences away. Therefore relative price levels should equal the exchange rate. The Figure 3 shows, however, that the guilder was severely undervalued in the 18th century, while the converse has been the case since the fifties of the last century.

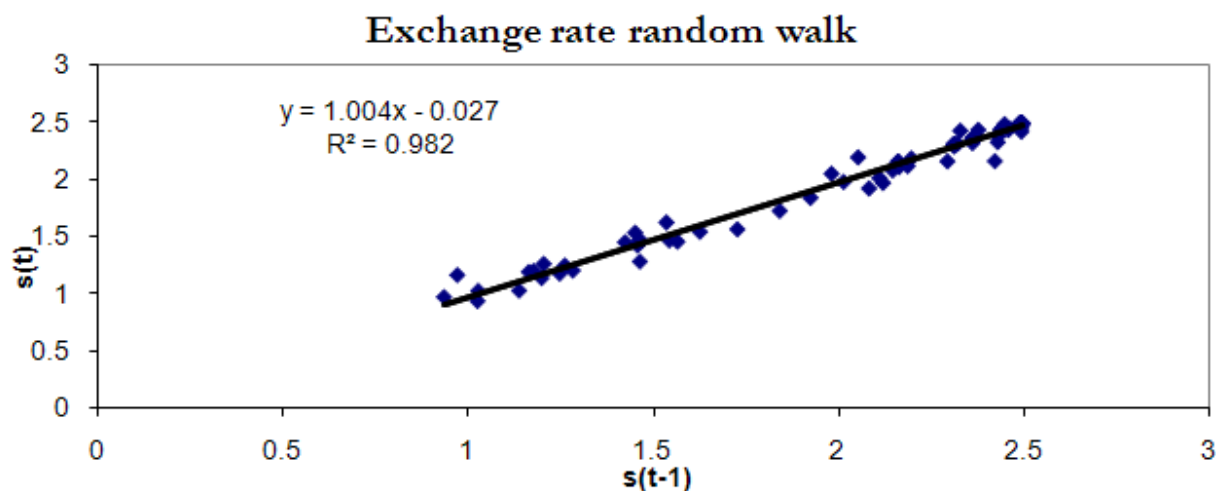


Until the middle of the 18th century, Dutch prices relative to English prices declined strongly. The exchange rate remained mostly stable due to the fact that both currencies used a metallic standard. This development in relative prices can be explained by the fact that The Netherlands lost a lot of its economic power during the last part of the 17th century and throughout most of the 18th century, while England prospered due to the industrial revolution and its hegemony at sea. Another important factor was the exorbitant import duties levied by England and France on Dutch exports.

3. Expectations

The lengthy undervaluation of the British pound during the post-war era can be explained by distrust towards English monetary policy. The exchange rate of fiat currencies is a weighted average of various macroeconomic variables on the one hand, and the expected value of the currencies on the other hand. When macroeconomic policies of two countries are the same, then deviations in the values of their currencies are simply background noise. The macroeconomic variables of one country relative to another country then follow approximately a random walk, meaning that their future values are best predicted by their current values, adjusted by the interest differential. In this case, the exchange rate depends solely on the relative values of macroeconomic variables. Alternatively, suppose the British authorities follow an expansionary policy which in the end needs to be financed in part by printing more money due to budgetary weakness of the government. Then the market expects inflation to pick up in the future. As a result, already today's exchange rate will be lowered, out of line with the current fundamentals, as today's rate must compensate for the expected future appreciation of the guilder.

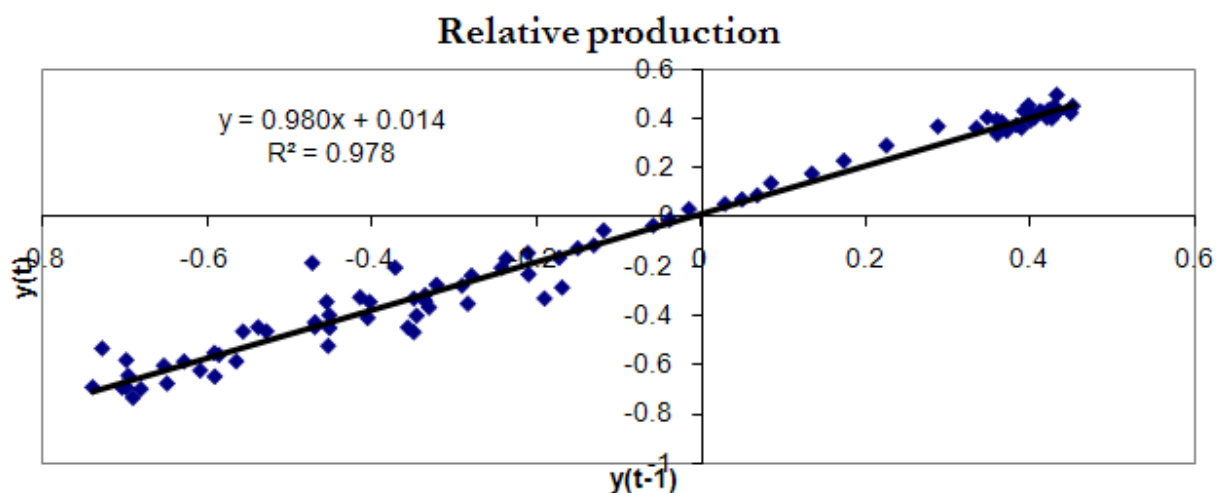

The next figure is a plot of the exchange rate of a given year cross plotted against the exchange rate in the previous year, starting from 1900.



If the exchange rate follows a random walk, then all points must lie on a line through the origin with a slope of 1. This appears approximately to be the case: A regression gives us a slope value of 1.004. The intercept however is slightly negative (-0.027) due to the prolonged fall of the pound. This result is precisely what the efficient market hypothesis predicts: All available information has been processed and is part of the current exchange rate, which means that any change in the rate, i.e. the difference between today's and yesterday's exchange rate, is caused entirely by newly available information.

Many economists have tried to predict the future values of exchange rates using variables such as the money supply and GDP. Of course, the exchange rate's random walk implies that its driving economic forces will also follow a random walk path, which means that macroeconomic indicators won't get you any closer to predicting exchange rates. This is evidenced by the following two figures, that cross-plot the money supply and relative production against their lagged values. We see again that the data is clustered around a 45-degree line with an intercept close to the origin. The money supply plot has a slope of 1.013 and an intercept which is again slightly negative at -0.018. Thus the relative fundamentals are as unpredictable as the exchange rate.


The random walk feature in fundamentals and exchange rates does not imply that the government can do whatever it wants. It does say that markets are capable of forming relatively accurate predictions of economic developments, both of exchange rates and fundamentals.



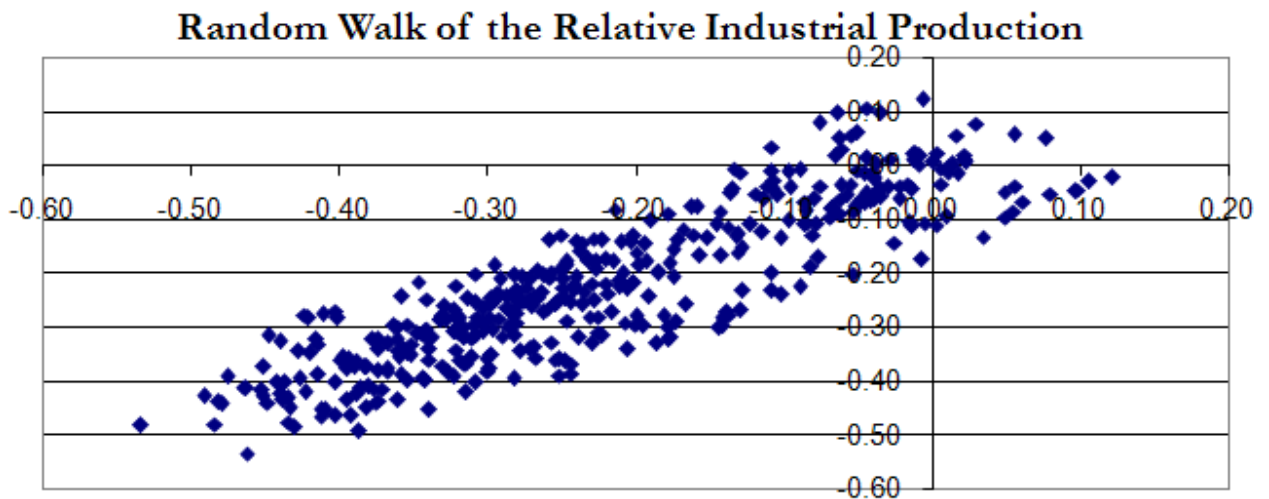

4. An Intermezzo

Here is another illustration of the random walk phenomenon, this time using the DM/Euro-US\$ rates. We start again by cross plotting the logarithm of the exchange rate $s(t)$ against its own past $s(t-1)$ for the monthly series. On purpose we do not correct for the change in denomination at the start of the fixing in 1999. This brings out clearly the period of the euro in the South-West part of the diagram. The market efficiency is again reflected by the fact that the data cluster around the diagonal. If the slope would be different, one could in effect forecast the direction of the future exchange rate. Currency markets would quickly arbitrage away such knowledge by taking positions today, net of the interest differential, thereby changing the current exchange rate in the direction of the predicted rate. As before, the plot reveals that, net of interest rates, the best forecast for outsiders is the current level.

Random Walk of the DM/EURO-US\$ Exchange Rate, Unadjusted



The canonical exchange model is the monetary model of the exchange rate. This model holds that the current exchange rate is determined by the relative money supplies, relative output levels and the interest differential. Relative meaning US levels divided by German levels. If the model has any bite, then one should expect to see for the explanatory fundamental variables the same pattern as for the exchange rate. In the graphs below we have cross plotted the logarithms of the (relative) fundamentals against their own past. For the money stock we take ante 1999 the German levels, and post 1999 we use the EMU levels. This shows the break in the unadjusted exchange rate cross plot is related to the introduction of the euro. The monthly industrial production series is for Germany only and hence shows no gap. The interest differential should not be affected by the introduction of the euro either, given that the ECB is essentially following in the footsteps of the Bundesbank.



Interestingly, all fundamental series nicely align around the diagonal. The random walk character of the exchange rate, producing the no change forecast, is thus also present in the fundamental series. One can see how the two views of the exchange rate can be made mutually compatible. Both the exchange and the fundamentals are equally difficult to predict and hence give the same result.

5.Relationship


Back to our Pound-guilder dataset, there *is* a relationship between the exchange rate and the relative macroeconomic variables, random walk or not. To show this, we regress the relative price levels against their own past values. Once again we find an almost perfect random walk

with a slightly negative intercept (-0,012) caused by the bouts of relatively high inflation in England.


When a variable is (covariance-) *stationary*, its variance and mean are fixed, and in the absence of shocks it will eventually return to its mean. A random walk is non-stationary per definition; it can reach arbitrarily high values. However, a combination of two random walk variables may very well be stationary; this relationship is called *cointegration*. If we examine the difference between the relative prices and the exchange rate, i.e. the real exchange rate, we find a slope coefficient of 0.94. This slope is significantly lower than 1. It implies that a shock to the real exchange rate has a half-life period of twelve years. It does take some time, but the exchange rates and relative prices are sufficiently interconnected, we say cointegrated, to effectuate a return to equilibrium.

6. Volatility


Any investor trades of risk against expected return. In the above we discussed the expected return or the mean of the exchange rate. It is time to move on to discuss the risk of foreign exchange investments. The variance or volatility is a popular measure of risk. The next figures show plots of the squared foreign exchange rate returns as a measure of volatility. We see that some periods clearly stand out from the rest. This is the well-observed clustering of volatile episodes and periods of quiescence that was first observed by Mandelbrot back in 1963. Later, Nobel Prize winner Robert Engle proposed an attractive autoregressive scheme, dubbed ARCH, which could exhibit the predictability in volatility while retaining the zero mean for the expected returns to reflect the absence of arbitrage opportunities. The ARCH scheme is now widely used by practitioners (such as option traders) to hedge their positions.




Monthly Volatility 1776-2000



Weekly Volatility 1915-2000




Interestingly, the clustering of low and high volatility is present in the data, no matter at what level of detail we choose for examining the data. Consider the plots for the monthly returns from 1766-2008, and the weekly returns from 1915-2008. If we compare these two figures with the previous plot for the yearly volatility, we see a very similar picture after adjustment for the scale. The bouts of volatility come in clusters regardless of the frequency of observations. Moreover, two or three outliers dominate each graph and determine the scale of the vertical axis. The remarkable observation we take away from these three figures is that volatility is self-scaling. This is nowadays known as the fractal nature of volatility. In Dutch scientific slang, the fractal property of speculative returns is referred to as the Droste effect, reminiscent of the [nurse](#) carrying a cup of hot chocolate and a can of Droste cocoa displaying a label with a picture of that same nurse, ad infinitum. No matter how much we zoom in or out, the same pattern in the data emerges.




7. Fat tails

Let us delve a bit deeper in the fractal nature of the exchange rate and asset prices in general. Before we can understand this self-scaling, we have to discuss a property of the unconditional return data. In almost all sciences, the normal distribution is widely used because averages have a tendency to be normally distributed. Furthermore, the normal distribution is the only finite variance distribution that is self-scaling. It is thus the obvious candidate for modeling financial returns. Below are histograms of the daily returns since 1971 and 1609, with a bell-shaped normal density curve overlay, estimated from the means and variances in the sample.



Daily returns



Yearly returns

It is clear to the naked eye that the normal model fails miserably. Relative to the normal model, the empirical density is too peaked, and has tails that are far too fat. We have to search for another candidate distribution. Mandelbrot (1963) initially proposed a model that would preserve both the self-scaling nature and the fat-tail feature, but the infinite variance assumption that comes with this model imposes too much tail fatness and does not generate the observed clustering of volatility. The three data features can be reconciled, however, if we drop the self-scaling requirement for all the data. It turns out that the clustering and the fat-tail property are commensurate with the self-scaling, if the latter property only applies in the tail area. This brings us to the work of the Italian engineer who took over Walras's chair at the university of Lausanne.

8. Pareto

Motivated by the social debates of his time, Pareto in 1896 discovered a remarkable time invariance regarding the distribution of the highest incomes. Due to this invariance Pareto concluded that the poor could be helped by a strategy for growth, but that redistribution would not work. This advice was hotly debated by many economists. Pareto's law survives today as the well-known Pareto distribution in statistics. Suppose we adapt Pareto's law to describe the distribution of the logarithmic loss returns X :

$$\Pr\{X \leq -x\} = ax^{-\alpha} \text{ for } \alpha > 0 \text{ and } x > a^{1/\alpha}$$

Note that the x represents losses, and


$$\Pr\{X \leq -x\}$$

stands for "the probability that the return X will be less than $-x$ " (i.e. a loss greater than x). One can understand why these distributions are fat tailed, since integration shows those moments larger than α are unbounded due to the explosion of the argument as x tends to infinity. Note that if we take logarithms on both sides of this equation, we obtain a linear relationship between the log-returns and their log-frequency with slope $-\alpha$:


$$\ln \Pr\{X \leq -x\} = \ln a - \alpha \ln x$$

If the same two sided logarithmic transformation is applied to an exponential type tail, such as the case for the normal distribution, one obtains a slope that is constantly changing, i.e. a slope that depends on the value of x . This is necessary for the boundedness of all moments.

What do the exchange rate data tell us? We first plot all the monthly loss returns logarithmically transformed against their log-rank order, when ranked in descending fashion. The second plot zooms in on the 50 highest losses only. If Pareto's model applies, one should see a straight line.



Pareto plot (monthly, loss returns)



Pareto plot (monthly, loss returns)

The highest losses in the second diagram indeed give more or less a straight line. Linear regression gives a slope coefficient of -2.26. This implies that moments up to the variance do exist, but the skewness (third moment) may very well be unbounded.

9. Synthesis

The three asset return data features can now be brought into agreement with each other if we assume that the tail of the distribution behaves like the Pareto distribution. This is similar to the case of the income distribution, where the Pareto model is known to fit well only for the highest income brackets. First, there are many fat-tailed distributions, such as the Student-t distributions, which exhibit the Pareto law as the first term of a (Taylor) expansion of the distribution in the tail area, but which are quite different in the centre. Second, our theoretical research has shown that the distribution of the extremes from an ARCH process, which does model the clustering of volatilities, indeed adheres to the Pareto-in-the-tail model. Third, a famous result in probability theory by Feller holds that if a distribution satisfies the Pareto model in the tail area, then for large x the sum of two independent draws from such a distribution satisfies the following equation:


$$\Pr\{X_1 + X_1 \leq -x\} \approx 2ax^{-\alpha}$$

Compare this result with the first equation above. Except for the factor two, the right-hand sides of both expressions are equal, implying self-scaling (with factor $z^{-1/\alpha}$). Apparently, one can just add the probabilities of the single losses to obtain the probability of the tail loss on the portfolio. Recall that log-returns are additive (i.e. the yearly return is just the sum of the monthly returns), and the relevance of the scaling feature becomes evident. Only recently has it been shown that this scaling also holds if the draws come from a process like ARCH, when the returns are time dependent. Nowadays, the Pareto-in-the-tail model is often used in risk management systems employed by the larger commercial banks to calculate their risk exposure to extreme events. The self-scaling feature is exploited to reduce the computational burden, since the exposure has to be computed for different investment horizons. By using the scaling law, it suffices to estimate the risk at the highest frequency and to extrapolate to lower frequencies by means of the scaling law.

Why do economic data exhibit the fat-tail property? Suppose, for example, that the asset price is uniformly distributed on the unit interval. It follows that the gross return, which is just the ratio of two consecutive prices, has a fat-tailed distribution (with $\alpha = 1$). The reason is that zero is in the support of the uniform distribution. The lagged asset price is in the denominator of the return. Hence, there is a positive probability that one has to divide by a very small number. This induces the fat tail nature of the returns. While this explanation is clearly far too simplistic, current research is actively addressing this question with more elaborate models.



10. Dependency

The fat tail property is not only the driving force between single risks, but also leaves its imprint on joint risks or systemic risk. To see this, we cross plot two different return series against each other. We choose two exchange rates against the Mexican peso (MXN): The Chilean peso (CLP) and the Columbian peso (COP). The next figure provides the cross-plot of the CLP/MXN versus the COP/MXN daily foreign exchange rate returns from 1994 to 2006.



To put this cross-plot into perspective, we produce two remakes of this plot under different distributional assumptions. The first remake assumes that the data are jointly normally distributed. It imposes the same means, variances and correlation (0.819) as are in the real data. Note that the normal remake fails dramatically as it does not produce the large outliers that we see in the real data. This is due to the exponential type tail, whereas the true data must be fat

tailed. Moreover, the remake does not display the strong clustering of the most extreme observations along the diagonal as we see in the real data. The normal cross-plot is also too elliptical. Hence, the normal also does not reveal the tail type of tail dependence that is in the data. Alternatively, we used the Student-t distribution with 3 degrees of freedom, see the last figure. The Student-t has Pareto type tails and also displays stronger tail dependence than the normal distribution. This latter remake does generate more outliers that, moreover, are closer aligned along the diagonal.



11.Explanation

Two basic conditions are sufficient for the frequent occurrence of systemic (widespread) market crises. First, each exchange rate distribution must have the heavy tail feature, meaning that the probability of a currency collapse is approximately Pareto distributed. Second, the (logarithm of) nominal bilateral exchange rates, expressed against the same base currency, are linear expressions of the domestic and the base currency fundamentals. In other words, the

fundamentals of the base currency partly drive the risk of both exchange rates. The standard monetary model of the foreign exchange rate provides such an affine framework. The two conditions together imply that joint currency crises will occur frequently and with vehemence.

The explanation applies to any group of assets whose values are linearly driven by a common set of risk factors that exhibit similar Pareto type tail risk. The common factor(s) overwhelm the idiosyncratic risk factors. Thus in the case of the two Latin American exchange rates, the common factors are the Mexican fundamentals. One can use the Feller result to show that the fundamentals of Chile and Columbia are dominated by the common Mexican factor. Since this factor is heavy tailed, it produces the outliers along the diagonal. Per contrast, if the factors were normally distributed, the common Mexican factor would induce positive correlation, but the dependency would quickly evaporate in the tails due to the exponential decline.