

Medieval foreign exchange: a time series analysis

Book or Report Section

Accepted Version

Bell, A., Brooks, C. and Moore, T. (2013) Medieval foreign exchange: a time series analysis. In: Casson, M. and Hashimzade, N. (eds.) Large Databases in Economic History: Research Methods and Case Studies. Routledge Explorations in Economic History. Routledge, Abingdon, pp. 97-123. ISBN 9780415820684 Available at <https://centaur.reading.ac.uk/35341/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Published version at: <http://www.routledge.com/books/details/9780415820684/>

Publisher: Routledge

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

This is the Authors' Accepted Manuscript. The definitive version has been published as Adrian R. Bell, Chris Brooks and Tony K. Moore, 'Medieval Foreign Exchange: A Time Series Analysis', in *Large Databases in Economic History: Research Methods and Case Studies*, ed. M. Casson and N. Hashimzade (Routledge, 2013), pp.97-123.

Medieval Foreign Exchange: A Time Series Analysis

Adrian R. Bell, Chris Brooks and Tony K. Moore

4.1. Introduction

This chapter will demonstrate some of the potential historical applications of econometric testing, focusing on a time series analysis of medieval foreign exchange rates. The work is part of a wider research project *Medieval Foreign Exchange, 1300-1500*, funded by the Leverhulme Trust under grant RPG193. The project is currently collecting a new database of medieval exchange rates and the following analysis is an exploratory study based on existing datasets intended to illustrate relevant forms of analysis. The first part of the chapter will briefly discuss the surviving evidence for medieval exchange rates. This is relatively extensive, at least by medieval standards, but there are a number of important caveats that need to be considered carefully. Despite this, exchange rate data offers some of the best opportunities to apply econometric analyses to the medieval economy. Recent studies have looked at volatility (Booth and Umit, 2008), market integration (Volckart and Wolf, 2006; Kugler, 2011; Chilosì and Volckart, 2011) and implicit interest rates (Booth, 2011). With the exception of some early work by Hyman Sardy (de Roover, 1968), however, economic historians have not really used modern time series analysis to study medieval exchange rates. The second part of the chapter will use such techniques to test for seasonality, stationarity and structural breaks. Finally, it will compare the results of these tests against work on modern exchange rates as well as contemporary evidence from medieval merchants' manuals.

4.2. Sources in their historical context

Reliable quantitative data on the medieval economy are hard to find, including key economic indicators. For example, there is very little direct evidence for interest rates (Bell *et al.*, 2009) or about trade and financial flows. Other information is more readily available – for instance, agricultural prices and wages can be extracted from manorial accounts (see Chapter 2), while final concords (fines) provide data about property values, as discussed in the following chapter. Another exception is foreign exchange rates (Einzig, 1970). This can be attributed to the confluence of three factors. The first was a general increase in long-distance trade following the 'commercial revolution of the thirteenth century' (Spufford, 1991). The second was the concurrent establishment of independent territorial states across Europe. Minting coins was both a key attribute of sovereignty and a source of income and so rulers sought to enforce the use of their local currency within their borders (Munro, 1979). As a result, merchants conducting long-distance trade needed to engage in foreign exchange. The third was the increasing sophistication of record keeping. In particular, the accurate recording and calculation of exchange rates was vital for merchants in order to keep their books balanced (de Roover, 1944a). While the vast majority of such records have been lost, enough survive to shed considerable light on medieval exchange rates.

Unfortunately, the surviving information about foreign exchange is scattered through a wide range of different sources. The most comprehensive dataset of medieval exchange rates was compiled by Peter Spufford (1986) in the 1970s and 1980s from correspondence with scholars across Europe. The *Handbook of Medieval Exchange* covers the period from 1106 to 1510, although the bulk of the data comes from the fourteenth and fifteenth centuries. It includes 13,197 exchange rate observations for 696 different currency pairs extracted from

528 different sources.¹ Unlike the more recent *Handbook of World Exchange Rates, 1590-1914* (Denzel, 2010), Spufford did not attempt to calculate annual average figures for each currency, as the extant data was not amenable to such treatment. Instead he presented individual raw exchange rates as found in the sources.

The aim of the *Handbook* was to provide indicative information about the relative values of different currencies that historians could use when drawing comparisons between countries. This meant that Spufford sought to cover as many currency pairs as possible over as long a period as possible. This inclusive approach necessarily generated a great variety in the nature of the evidence used. At the same time, there were also limitations imposed by the format of the *Handbook* as a printed book. Instead of listing all cross-exchanges between all currencies, Spufford concentrated on listing the exchange rates of each currency against the major international currency; first, the Florentine florin and, from the fifteenth century, the Venetian ducat. This is similar to the use of the dollar today. Moreover, even where more frequent data were available, Spufford only selected a maximum of one rate for each month.

As a result, the *Handbook* presents a number of problems for the purpose of quantitative analysis, although these can better be seen as the flip-side of its strengths. In particular, the *Handbook's* inclusive approach to data collection, necessary given the scattered nature of the surviving source material, means that the data are not always strictly comparable. Spufford himself stresses this fact and one of the great strengths of the *Handbook* is that it provides information about the type of transaction, its location and, vitally, a reference to the original source, so that the exchange rates can be checked.

Most obviously, the *Handbook* includes exchange rates taken from a variety of different types of source and relating to different types of exchange transaction (Spufford, 1986, pp.l-liii). For example, spot exchange of one coin for another did not involve a time element whereas bills of exchange did, as will be shown below. The exchange rates used in accounts could also become fossilised at anachronistic rates or could be manipulated, for example to hide interest charges, while the official exchange rates promulgated by governments could be aspirational rather than effective. As a result, the exchange rates quoted in different types of transactions could vary considerably, making it difficult to draw valid comparisons.

Another important factor influencing medieval exchange rates, and a major difference to the modern situation, was the location of the transaction. In part, this reflected the slower speed of medieval communications. It also resulted from patterns of trade, which affected the demand for and supply of money differently at different places. More fundamentally, the bill of exchange, the classic foreign exchange instrument, did not just include a foreign exchange transaction but necessarily involved the extension of credit, since the buyer of the bill paid the seller upfront in local currency but only received the value in foreign currency later. Bills of exchange between different financial centres had different maturities (known as *usance*), mostly increasing with geographical distance. The exchange rates quoted at different financial centres therefore incorporated a spread to account for the time value of money. The operation of this system was demonstrated by de Roover (1944b) and has been traced back to the later thirteenth century by Thomas Blomquist (1990, pp.362-8). It has even been argued that the primary significance of the bill of exchange was that it enabled merchants to circumvent the usury prohibition on charging interest (Rubin, 2010; Koyama, 2010). Other

¹ The data from the *Handbook of Medieval Exchange* has been uploaded to the Medieval and Early Modern Data Bank and can be accessed at http://www2.scc.rutgers.edu/memdb/search_form_spuf.php.

historians have argued that this over-states the credit aspect of the bill of exchange and downplays its role in facilitating international trade (Leone, 1983). The important point for our current purposes is that exchange rates varied depending on the location (or the direction) of the transfer and so rates at different places cannot be compared directly.

Furthermore, where exchange rates have been taken from records of individual transactions, there may be idiosyncratic factors influencing the rates quoted in that particular case. To take one example, in May and June 1305 the London branch of the Gallerani of Siena sold seven bills of exchange to customers wishing to remit money from England to the papal curia (Bigwood and Grunzweig, 1962). The rates quoted by the Gallerani varied from 5¾ florins per mark sterling on 21 May to 4½ florins per mark sterling on 7 June, a drop of 21.7 per cent in a little over two weeks. However, by 14 June the exchange rate had rebounded back up to 5¾ florins per mark sterling, an increase of 27.8 per cent in one week. It is unlikely that market exchange rates really fluctuated this wildly. A more plausible explanation is that the Gallerani offered different exchange rates to different customers. For instance, the three buyers that received the most favourable rates were all connected to the papacy, while the worst rate was received by an English clerk.

Another fundamental challenge concerns the frequency and distribution of the surviving evidence. Even for the best documented currency pairs, the *Handbook* records an average of only 1-2 observations per year. This means that idiosyncratic factors influencing individual exchange rates could potentially distort long-run trends. Also, there were strong seasonal variations in exchange rates, as will be shown below. Since the variation in exchange rates within each year was usually greater than the change from year to year, long run developments might be obscured if the data for some years quoted rates from a seasonal peak while that for others came from seasonal lows. There are frequent gaps in the series, which raises problems of interpolation. Finally, the frequency of observations varies dramatically over time, which makes it difficult to conduct many types of statistical analysis.

In this chapter, we take a different approach to constructing a dataset of medieval exchange rates. Rather than seeking to cover a long time span, we focus on a relatively short period for which we have abundant data. Instead of following an inclusive data collection policy, our approach is more exclusive and limited to rates from one particular type of source, namely mercantile correspondence. The merchant who had better and more up-to-date information about exchange rates in other banking centres enjoyed an advantage over his uninformed peers. For this reason, merchants often listed the current market exchange rates at the end of their commercial letters. These rates were probably collected from the bill brokers that arranged deals in each city (de Roover, 1968, p.29). They can be seen as forerunners of the exchange rate currents printed from the sixteenth century onwards (McCuster and Gravesteijn, 1991), which ultimately developed into the modern financial press. The use of exchange rates from commercial correspondence has two main advantages. First, the market rates stripped out some of the idiosyncratic factors that may have influenced the exchange rates used in particular transactions. This makes the data more useful for comparative purposes, although historians should bear in mind Mueller's (1997, p. 588) warning against 'fetishiz[ing]' the rates quoted in mercantile correspondence since the rates used in actual transactions may have varied depending on the relative 'contractual leverage' of the two parties. Second, merchants wrote frequently to their correspondents (on a weekly or even daily basis) and often corresponded with several different cities, providing a greater depth and higher frequency of data.

This chapter draws upon the archive of one merchant in particular, namely Francesco di Marco Datini of Prato, near Florence in Italy. Datini was an extremely successful merchant, and between 1380 and 1410 his network of branches and correspondents covered much of western Europe (Origo, 1963; Nigro, 2010). The contents of the archive are described in Federigo Melis (1962) and many of the documents have been digitised and can be consulted online.² The potential of the Datini archive has long been recognised, and the commercial letters have been mined for data about exchange rates; de Roover (1968) extracted the exchange rates cited in letters from Barcelona and from Bruges while Reinhold Mueller (1997) did the same for the letters from Venice.³ The *Medieval Foreign Exchange* project is currently extending this dataset to cover Florence and Genoa, the other two key banking centres in medieval Italy. The enlarged dataset has not been finalised, however, so the following statistical analysis is based on the data collected by de Roover and Mueller.

4.3. Descriptive statistics

Time series analysis is a cornerstone of econometrics and the evaluation of a series' properties when viewed on its own should be a precursor to any solid multivariate research. We will first describe the data, in particular focusing on the statistical features of the series; we then move on to discuss the exchange rates' seasonal patterns and outline a framework for determining whether the series are best described as stationary or non-stationary processes. Finally, we consider in detail, using two different approaches, how to determine whether there were structural breaks.

As explained above, the following analysis draws on exchange rates quoted in merchants' letters from Barcelona, Bruges and Venice written between c.1385 and c.1410. The basic sources are shown in Table 4.1. The first two columns list the currency pair and the location. This effectively shows the direction of the exchange (e.g. Barcelona-Bruges is the exchange rate at Barcelona between the lira of Barcelona and the Flemish écu while Bruges-Barcelona is the rate at Bruges for the same pair). The third column describes how the exchange rate was quoted. The following analysis is based on monthly percentage changes in the exchange rates rather than raw levels but the method of quotation is still very significant for understanding the significance of the growth rates of the exchange rates discussed below. For example, Bruges 'gave certain' to Barcelona and London; that is, the exchange rate was quoted as an uncertain number of pence of Barcelona or pence sterling per écu of Bruges. Thus, a rise in the quoted exchange rate reflects an increase in the value of the écu. Bruges, however, 'gave uncertain' to Genoa, Pisa, Paris and Venice, meaning that exchange rates were quoted as a variable number of Flemish groats per unit of foreign currency. Thus an increase in these exchange rates actually means that the écu was declining in value.

[TABLE 4.1 HERE]

² The homepage of the Datini archive is <http://datini.archiviodistato.prato.it/www/indice.html>.

³ The exchange rates for Barcelona and Bruges were entered into a spreadsheet from the appendix in de Roover (1968). Mueller's data for Venice is available from the Medieval and Early Modern Data Bank (http://www2.scc.rutgers.edu/memdb/search_form_mueller.php). It should be noted that there are a number of errors in the online dataset, especially for the Venice-Barcelona exchange rates. These do not seem to be included in Mueller's own figures and presumably occurred during when the data was being re-entered for upload. We identified data points that deviated significantly from contemporaneous rates, checked them against the original letters from the Datini archive and made corrections where necessary.

The penultimate two columns show the total number of data points available for analysis and the main date range over which we have observations. We take monthly averages over all data points available for that month. It is clear that with the notable exceptions of Barcelona-Venice, Bruges-Pisa, and Venice-Rome, where the samples are very small, for most of the series we have at least twenty years of monthly data or typically 200-300 observations. Unfortunately, it is in the nature of medieval data sources that there are inevitably missing values for some months. In order not to lose too many data points, we interpolate in such cases by rolling forward the value that was available for the previous month. The final column of Table 4.1 shows the number of interpolated values for each series. As can be seen, the number of such missing values is modest, and never more than 10 per cent of the total sample. We ensure that we never roll forward a data point for more than three months in a row – if there are more than three months’ of missing observations, we truncate the sample at that point.

We work mainly with the growth rates of the exchange rates, since we know that this analysis will be econometrically valid even if the raw rates contain unit roots, as discussed below. Thus we start by presenting the main features of the exchange rates in Table 4.2. The table presents summary statistics for the monthly percentage changes of the 25 series that we examine spanning the three main venues of Barcelona, Bruges and Venice.

[TABLE 4.2. HERE]

It is evident that the Venetian means are all positive (except for Venice-Paris where Venice ‘gave uncertain’ to Paris), indicating the strength of the ducat over the period. In particular, the ducat rose by around 0.2 per cent per month (2.4 per cent per annum) against the Roman florin and by 0.14 per cent per month (1.6 per cent per annum) against the pound sterling. This reflects the maintenance of the gold content of the ducat, as well its increasing displacement of the Florentine florin as the major international coin (Spufford, 1991, p.321).

The variance estimates, presented in the fourth column, are broadly similar to modern figures, which would be of the order of 3-5 per cent per month. The variance, sometimes known in finance as the volatility of a series, measures the extent to which it moves around over time. It thus shows the spread of the observations around their mean value. However, we can note considerable differences in volatilities across the series – from just 0.8 per cent for Barcelona-Majorca to 9.5 per cent for Venice-Bruges and 13 per cent for Barcelona-Bruges. The former probably reflects the fact that Barcelona and Majorca were both part of the Aragonese realm. The latter are almost certainly caused by outliers in the percentage change series for these currency pairs as a result of currency revaluations in Bruges (Munro, 2012).

In terms of their symmetry or otherwise, some of the series are left-skewed and some are right-skewed; all series are leptokurtic, implying that they have more in the mean, fatter tails but less in the shoulders of the distribution compared with a normally distributed series having the same mean and variance. Again, this feature is very similar to that in standard contemporary asset return series. This leptokurtosis is a key feature of financial time series from a risk management perspective since it implies that extreme movements (i.e. very large swings) are more likely than would be the case under a normal distribution. Barcelona-Bruges is the most skewed and has the highest kurtosis, in particular due to an almost 50 per cent rise between October and November 1390. This reflects a curious delayed impact of the enhancement of the Flemish currency in 1389. The Italian merchants at Barcelona continued

to quote the exchange rate with Bruges in terms of the old money until November 1390, when they switched to using the new money (de Roover, 1968, pp.39-40).

4.4. Seasonal patterns in exchange rates

Since bills of exchange were used both to transfer money for trade and also to borrow or lend money, exchange rates were closely linked to the wider condition of the money market. The sixteenth-century merchant Bernardo Davanzati explained this relationship using the analogy of a hand tightening or loosening its grip on money (Mueller, 1997, p.305). At times of high demand for cash, the hand would tighten and not release any money except at a higher price. As a result, the exchange rate would rise, that is, sellers of bills of exchange (borrowers) would have to promise more foreign currency to receive one unit of local currency. In this case, merchants described money as being ‘dear’ (*carestia*) or ‘tight’ (*strettazza*). In the contrary situation, when the supply of money exceeded the demand, money was ‘loose’ (*larghezza*) or ‘abundant’ (*dovizia*) and exchange rates fell as buyers of bills of exchange (lenders) were prepared to accept less foreign currency per unit of local currency. According to Giovanni di Antonio da Uzzano, the ‘good rule in making exchange’ was to anticipate changes in exchange rates and not merely to react to them (Pagnini, 1766, p.153). It is therefore important to know whether there were predictable seasonal patterns in medieval exchange rates.

There are various ways to test for seasonality in time-series data. It is possible to employ trigonometric functions or to work in the frequency domain. However, in such cases the quantitative sophistication arises at the expense of interpretability. A much simpler approach, which we apply here, is to use a linear regression including monthly dummy variables

$$y_t = \alpha_1 D1_t + \alpha_2 D2_t + \alpha_3 D3_t + \alpha_4 D4_t + \alpha_5 D5_t + \alpha_6 D6_t + \alpha_7 D7_t + \alpha_8 D8_t + \alpha_9 D9_t + \alpha_{10} D10_t + \alpha_{11} D11_t + \alpha_{12} D12_t + u_t \quad (.1)$$

where y_t is the exchange rate percentage change series under consideration, and $D1_t, D2_t \dots D12_t$ are monthly dummy variables for January, February, ..., December, and u_t is an error term, assumed to be normally distributed with zero mean and constant variance. These dummy variables take the value one for the month to which they correspond and zero otherwise – so, for example, $D1_t$ takes the value one every January and zero for every other month. This way, the dummies effectively “pull out” the observations for their corresponding months and set everything else to zero. Thus each parameter attached to the dummies can be interpreted as the average change in the foreign exchange rate for that month (in percentage terms).

Note that, given the way that it has been specified to contain a full set of 12 monthly dummy variables, this regression must not contain an intercept term to avoid the “dummy variable trap.” This would have arisen if all possible seasonal dummies given the frequency of the data employed (i.e. 12 for monthly, four for quarterly, etc.) were included in the model together with an intercept. The result would be that the regression could not be run due to a multicollinearity (see Chapter 1).

The seasonality results are presented in Table 4.3 for the rates quoted in Barcelona, in Table 4.4 for the rates quoted in Bruges, and in Table 4.5 for those quoted in Venice. Overall, these tables indicate very clear patterns. This quantitative evidence can be used to test the qualitative descriptions of the state of the money market by contemporary merchants. In this

case, we shall use Giovanni di Antonio da Uzzano's *Practica Della Mercatura* (Pagnini, 1766). It should be noted that Uzzano was writing c.1442, roughly fifty years after our data from the Datini letters, and this may explain some of the discrepancies between our statistical reconstruction and Uzzano's experience.

[TABLES 4.3, 4.4. AND 4.5 HERE]

The foreign exchange rates quoted in Barcelona rise in December and most significantly in January and this continues at a more modest pace until around April, followed by large reversals in June through to October, with the most significant falls in the autumn months. In all cases, these exchange rates were quoted as an uncertain number of pence of Barcelona for a certain number of foreign coins. Thus a rise in exchange rates at Barcelona meant that the lira was decreasing in value, reflecting an easing of the money market. Conversely, a fall in exchange rates indicates an increase in the value of the *lira* and thus a tightening of the money market. This agrees with the first part of the account given by Uzzano:

In Barcelona, money is dear from the first of June through all of August because of the investments in wool from Aragon and the surrounding valleys and because of the purchases of 'grain' [the dye] in Valencia; the money market tightens again in October, after St Luke's day, which is on the 18th, because of the investments in saffron, when the dearness is even greater than in the wool season, and it will last until January; and from then on, money eases every day and the exchange rates return to their former level, and the easiness lasts until the wool season, unless something unexpected happens (Pagnini, 1766, p.156; translated in de Roover, 1968, p.88).

The Datini data support the tightness in the summer and autumn months (June to October) but the market seems to ease in the winter, when Uzzano suggests that the market should have continued to tighten and to peak. It is possible that the Spanish saffron trade was less important during the Datini period than it later became, although Datini himself traded in saffron (Origo, 1963, p.97).

For the foreign exchange rates quoted in Bruges, the pattern is less clear-cut, with most of the changes not being statistically significant. In general, we see that the exchange rates for Bruges-Barcelona and Bruges-London fell in the winter months from December to March (May for London) and rose from April (June for London) until November, while the other rates fell in the summer and rose in the winter and spring. Here it is important to recall that Bruges 'gave certain' to both Barcelona and London and a rise in these exchange rates therefore reflects a rising value for the Flemish *pond groot* and thus a tightening of the money market while a fall in the same rates reflects an easing of the market. However, the exchange rates with Genoa, Paris, Pisa and Venice were quoted as an uncertain number of Flemish groats for a certain number of foreign coins, and so the reverse logic applies. The evidence of the Datini letters indicates that the Bruges money market tightened in the summer and eased in the winter. Interestingly, this directly contradicts Uzzano's depiction of the seasonal trends in the Bruges' money market c.1442:

In Bruges, money is dear in December and January because of the many ships that are being loaded with commodities and dispatched at this time, and, in August and September, money expands because of the fairs that are being held and that attract merchants who come to purchase and bring in ready cash (Pagnini, 1766, p.156; translated in de Roover, 1968, p.90).

This discrepancy would repay further historical investigation into the seasonal patterns of trade in medieval Bruges. In particular, were there significant changes between the later fourteenth and early fifteenth centuries, when the Datini letters were written, and the mid-fifteenth century, when Uzzano wrote his manual, that might explain the stark differences set out above?

For the rates quoted in Venice, there are mostly significant falls in September and in February but rises in June and July. As Venice gave certain to most other currencies, this means that the market was tightening in the summer and loosening after September. We may note that Venice-Paris appears to show the opposite pattern, because here Venice 'gave uncertain' to Paris, but in fact shares the same underlying trend. This supports Uzzano's depiction of the Venetian money market :

In Venice, money is expensive from May to 8th September, because of the outward bound galleys which leave in July, August and September. The reason why it gets more expensive is because everyone starts to make arrangements and they want to remit more there; and this higher cost is due to the amount of cash the galleys carry, because a great deal of merchandise is sold there at the time of the galleys, which must be paid for just when you have many demands on your purse - and a lot of money goes out of the banks in cash, so cash is always dear there by 1 per cent more than usual. And money is highly priced for all places, and is offered there at various maturities. From 8th July money is highly priced, then there are no more maturities until 1st August, and in this month there is an expansion by ½ to 1 per cent. From 1st August money starts to fluctuate, and is expensive continually until 8th September; and after the 8th all maturities have become due, and all the galleys have gone, so there is no more demand - and the banks are quick to supply and money goes through the floor. (Pagnini, 1766, pp.156-7; translation by Dr Helen Bradley).

Mueller (1997, pp.305-7) adds further details. The galleys to Romania left Venice in mid-July, those to Beirut in mid-August, and finally the Alexandria galleys in late August to early September. Since western Europe had a trade deficit with the East, merchants had to export silver bullion to purchase goods for re-import. This led to a great demand for money in the summer, which drove up exchange rates. Subsequently, after the last galleys had departed, the demand for money reduced and thus exchange rates collapsed in September. Uzzano also mentions a later tightening of the money market in December and January, linked to the departure of galleys to Catalonia (Pagnini, 1766, p.157). This fits neatly with the sharp rise in the Venice-Barcelona exchange rate in December shown in Table 4.5.

A further important issue is whether we analyse the seasonal patterns in the levels of the series or in their percentage changes. Where a series in levels has little underlying trend, it is likely to make very little difference. For instance, our results in Table 4.5 for Venice-Florence are very similar to the patterns in Graph 8.1 of Mueller (1997, p.307) where the sample period is very close to ours and Table 4.2 indicates that over time the mean change in the exchange rate was only 0.01 per cent per month. On the other hand, the pattern that we observe for Barcelona-Bruges is quite different to that reported by Hyman Sardy (de Roover, 1968) and this may be a result of his use of levels on a series where the mean percentage increase in the exchange rate was 0.09 per cent per month, almost ten times higher than that for Venice-Florence. When a series is trending heavily over time – either up or down – then any statistical analysis using the levels will not have the meaning that the researcher probably

intended. In the case of a series trending upwards, an analysis of the seasonal patterns will give undue weight to the observations at the end, which by definition will be much larger in value than those close to the beginning. Thus we would recommend historians to use percentage changes rather than levels in such cases.

The R^2 figures from these regressions measure the degree to which the seasonal patterns can explain the overall variation in each series. In an earlier analysis of the same data from Bruges and Barcelona, Hyman Sardy found that ‘about ten per cent of the fluctuations in the series could be attributed to seasonality’ (de Roover, 1968, p.104). Our work suggests that seasonality played a greater role, accounting for around twenty per cent of the variation in most series. This is highest for Bruges-Pisa (37 per cent of the variation explained, although based on the smallest sample) and for Venice-Pisa (27 per cent explained). This suggests that Pisa was particularly affected by seasonal flows of trade or finance. Seasonality is less significant for Bruges-Paris (10 per cent) and Venice-Bologna (11 per cent). Uzzano explains that there was a particularly close connection between exchange rates at these and a number of other places (Pagnini, 1766, p.154), which Mueller describes as ‘paired cities’, since their geographical proximity meant that information could be transmitted between the two in 1-2 days, allowing merchants in one location to react quickly to changes at the other (Mueller, 1997, pp.588-9).

4.5. Testing for unit roots in exchange rates

A key early question in analysing time-series is whether each variable in a sample can be considered stationary or whether it contains a unit root, which is so-called because the root of a characteristic equation for such a process is unity. A stationary series is one with a constant mean, constant variance, and constant autocovariance structure for a given lag – in other words, the relationship between the current value of the time-series and its previous values remains constant. This is probably the most important characteristic of a series as it has the most significant impact on its properties and also on which type of analysis is most appropriate. There is also the issue that if a series contains a unit root, standard econometric approaches cannot be applied to the data in their raw, levels form, otherwise this would result in ‘spurious regression’ where entirely independent unit root processes appear to standard econometric approaches to be strongly related. Thus, in such cases, the series must be converted into a percentage changes (i.e. growth rates) form.

However, whether exchange rate series in levels are non-stationary is surprisingly still an open question, with much empirical evidence both for and against. In the “modern finance” literature, much research has already been undertaken in this regard, even when we focus specifically on exchange rates. Many authors use the augmented Dickey Fuller (ADF) test (Dickey and Fuller, 1979), described below, on data of monthly or quarterly frequency. Meese and Singleton (1982), for example, test for the presence of a unit root in the log of a number of weekly US dollar-denominated exchange rates from the 1970s. They cannot reject the null hypothesis of a unit root in the log of the levels, but the percentage changes are found to be stationary, a result echoed by Corbae and Ouliaris (1988) in their tests on monthly Sterling data.

More recently, Whitt (1992) compared the results of a test for unit roots on the real US dollar exchange rate against a number of others using the ADF approach and Sims’ Bayesian test. The ADF test cannot reject the non-stationary null, but the Bayesian test does strongly reject it. Some evidence for stationarity in real exchange rates was also found by Taylor (1990). In a

comparative study, Schotman and van Dijk (1991) use the ADF and Sargan Bhargava unit root tests, together with a Bayesian posterior odds approach. The ADF and Sargan-Bhargava tests both lead to non-rejection for all the exchange rates they test, while the Bayesian method does not provide conclusive results.

Finally, Goodhart *et al.* (1993) conduct ADF and Phillips-Perron tests on a number of US dollar-based exchange rates at sampling frequencies from tick-by-tick to daily. The Phillips-Perron procedure is very similar to the ADF test described below except that it incorporates an automatic correction for autocorrelated residuals in the test regression. Their main finding is that daily and hourly series have a unit root, while most minute-by-minute and tick-by-tick series have a unit root with trend. They conclude that temporal aggregation preserves the non-stationarity in exchange rates so that the series are non-stationary whether they are observed hourly, every minute or as the transactions occur.

If a series y_t is initially non-stationary, but becomes stationary on differencing d times, then it contains d unit roots and is said to be integrated of order d . We “first difference” a series by subtracting the immediately previous value of that series from the current one. It is then possible to second difference a series by applying the same process again to the series that has already been differenced once. A unit root is only one possible form of non-stationarity, an explosive root being the other case, although the latter is rarely considered since it is more difficult to justify from an economic theoretical perspective. Many economic series contain an exact unit root when transformed into logarithms (Banerjee *et al.*, 1993, p99). Thus if we find that the raw data contain a unit root, but the percentage changes are stationary, this provides an *ex post* justification for using the percentage changes in subsequent analysis.

This section will now give a brief description of the Dickey-Fuller approach to testing for a unit root in time series. They conducted the early and pioneering work on this topic, and despite numerous advances in the testing theory since then, their technique still constitutes the main workhorse of unit root testing – see Dickey and Fuller (1979) and Fuller (1976). The Dickey-Fuller (DF) tests used in this chapter are the based on a t -ratio from a test regression. It is possible to include either a constant, or a deterministic trend, or both or neither in the test regression. Consistent with the expected features of the data, we elect to adopt the model including an intercept but not a trend. The null (H_0) and alternative (H_1) models for the test we implement are thus

$$H_0: y_t = y_{t-1} + u_t \quad (4.2)$$

$$H_1: y_t = \varphi y_{t-1} + \mu + u_t, \quad \varphi < 1 \quad (4.3)$$

This is a test for a random walk against a stationary autoregressive model of order one (AR(1)) with a drift. We can thus write

$$\Delta y_t = u_t \quad (4.4)$$

under the null hypothesis, where $\Delta y_t = y_t - y_{t-1}$, and the alternative may be expressed as

$$\Delta y_t = \psi y_{t-1} + \mu + u_t. \quad (4.5)$$

The test for whether the series contains a unit root is based on the t -ratio of the y_{t-1} term in the estimated regression and thus the test statistic is defined as

$$\text{Unit root test statistic} = \frac{\hat{\psi}}{\hat{SE}(\hat{\psi})} \quad (4.6)$$

The test statistic does not follow the usual t -distribution under the null hypothesis, since the null is one of non-stationarity, but rather it follows a non-standard distribution. Critical values are derived from Monte Carlo experiments in, for example, Fuller (1976). The null hypothesis of a unit root is rejected in favour of the stationary alternative in each case if the test statistic is more negative than the critical value.

The test may also be “augmented” by the addition of p lags of the dependent variable to the estimated equation, known as the augmented Dickey-Fuller test, which allows for possible autocorrelation of residuals in the regression. The unit root test would not perform well if such a structure was present in the residuals of the test regression but unaccounted for. The model under the alternative in this case can now be written as

$$\Delta y_t = \mu + \psi y_{t-1} + \sum_{i=1}^p \alpha_i \Delta y_{t-i} + u_t \quad (4.7)$$

A problem now arises in determining the “optimal” number of lags of the dependent variable to add to the estimated equation in order to sufficiently allow for autocorrelation, but not to over-fit. One way to determine the number of lags is to use the frequency of the data as a decision rule (e.g., use four lags for quarterly data, 12 for monthly and so on), but this is likely to lead to considerable over-fitting in the case of monthly data as we have here. Alternatively, the number of lags may also be determined, based on the data, using some kind of information criterion, such as Akaike’s or Schwarz’s Bayesian criteria. However, given the limited number of data points at our disposal and for consistency and comparability across series, we employ an arbitrary zero and three lags for all currency pairs investigated. We find that the conclusions are not qualitatively affected by this choice of lag length.

4.6. Structural breaks in exchange rates

A structural break occurs when the properties of a series go through a substantial change in behaviour so that previous models that described the relationship between variables subsequently break down. Although the term is often used somewhat loosely in the modern applied econometrics literature, it may be helpful to distinguish between a structural break and an outlier. A structural break is as described in the previous line, where the behaviour of a series or its relationship with other series changes on a long-term or permanent basis, whereas an outlier occurs when the properties of a series change for one or perhaps several periods before fairly quickly reverting back to its previous behaviour. In the context of medieval exchange rates, structural breaks could be caused by wars, financial crises, currency shortages, debasements or recoinages, poor harvests, and so on.

In terms of the econometrics, early tests for structural breaks were conducted in the context of regressions based on stationary time-series – for example, the Chow (1960) and Quandt (1960) likelihood ratio tests. However, some more recent approaches have been conducted in the context of tests for unit roots since it has been shown that the standard Dickey-Fuller-type unit root tests presented above do not perform well if there are one or more structural breaks in the series under investigation, either in the intercept or the slope of the regression. More

specifically, the tests have low power in such circumstances and they fail to reject the unit root null hypothesis when it is incorrect as the slope parameter in the regression of Δy_t on y_{t-1} is biased towards unity by an unparameterised structural break. In general, the larger the break and the smaller the sample, the lower the power of the test. As Leybourne *et al.* (1998) have shown, in addition unit root tests are oversized in the presence of structural breaks, so they also reject the null hypothesis too frequently when it is correct.

Perron's (1989) work was the first to systematically address the issue of testing for unit roots in the presence of structural breaks. This work is considered important since he was able to demonstrate that if we allow for structural breaks in the testing framework, a whole raft of macroeconomic series that Nelson and Plosser (1982) had identified as non-stationary may turn out to be stationary. Perron argues that most economic time series are best characterised by broken trend stationary processes, where the data generating process is a deterministic trend but with a structural break around 1929 that permanently changed the levels (i.e. the intercepts) of the series.

Perron (1989) proposes three test equations differing dependent on the type of break that was thought to be present. The first he terms a 'crash' model that allows a break in the level (i.e. the intercept) of the series; the second is a 'changing growth' model that allows for a break in the growth rate (i.e. the slope) of the series; the final model allows for both types of break to occur at the same time, changing both the intercept and the slope of the trend. If we define the break point in the data as T_b , and D_t is a dummy variable defined as

$$D_t = \begin{cases} 0 & \text{if } t < T_b \\ 1 & \text{if } t \geq T_b \end{cases} \quad (4.8)$$

The general equation for the most general type of test Perron proposed is

$$\Delta y_t = \mu + \psi y_{t-1} + \beta_1 D_t + \beta_2 (t - T_b) D_t + \lambda t + \sum_{i=1}^p \alpha_i \Delta y_{t-i} + u_t \quad (4.9)$$

For the crash only model, set $\beta_2 = 0$, while for the changing growth only model, set $\beta_1 = 0$. In all three cases, there is a unit root with a structural break at T_b under the null hypothesis and a series that is a stationary process with a break under the alternative.

While Perron (1989) initiated a new literature on testing for unit roots in the presence of structural breaks, an important limitation of his approach is that it assumes the break date is known in advance and the test is constructed using this information. It is possible, and perhaps even likely, however, that the date will not be known and must be determined from the data. More seriously, Christiano (1992) has argued that the critical values employed with the test will presume the break date to be chosen exogenously and yet most researchers will select a break point based on an examination of the data and thus the asymptotic theory assumed will no longer hold.

As a result, Banerjee *et al.* (1992) and Zivot and Andrews (1992) introduce an approach to testing for unit roots in the presence of structural change that allows the break date to be selected endogenously. Their methods are based on recursive, rolling and sequential tests. For the recursive and rolling tests, Banerjee *et al.* propose four specifications. First, the standard

Dickey-Fuller test on the whole sample, which they term \hat{t}_{DF} ; second, the ADF test is conducted repeatedly on the sub-samples and the minimal DF statistic, \hat{t}_{DF}^{\min} , is obtained; third, the maximal DF statistic is obtained from the sub-samples, \hat{t}_{DF}^{\max} ; finally, the difference between the maximal and minimal statistics, $\hat{t}_{DF}^{diff} = \hat{t}_{DF}^{\max} - \hat{t}_{DF}^{\min}$, is taken. For the sequential test, the whole sample is used each time with the following regression being run:

$$\Delta y_t = \mu + \psi y_{t-1} + \alpha \tau_t(t_{used}) + \lambda t + \sum_{i=1}^p \alpha_i \Delta y_{t-i} + u_t, \quad (4.10)$$

where $t_{used} = T_b/T$. The test is run repeatedly for different values of T_b over as much of the data as possible (a 'trimmed sample') that excludes the first few and the last few observations (since it is not possible to reliably detect breaks there). Clearly, it is $\tau_t(t_{used})$ that allows for the break, which can either be in the level (where $\tau_t(t_{used}) = 1$ if $t > t_{used}$ and 0 otherwise); or the break can be in the deterministic trend $\tau_t(t_{used}) = t - t_{used}$ if $t > t_{used}$ and 0 otherwise). For each specification, a different set of critical values is required, and these can be found in Banerjee *et al.*

Much recent work on whether exchange rates contain unit roots has been conducted in the panel context, where the additional information from combining series together can lead to considerable efficiency gains and improvements in power. Relevant research includes Jorion and Sweeney (1996), who reject the unit root null hypothesis for ten US dollar-denominated currencies and seven Deutschmark denominated currencies. Wu (1996) concludes firmly that US dollar real exchange rates are not unit root processes using a panel of 18 monthly real exchange rate series; a similar result is observed by Oh (1996) and MacDonald (1996) using longer samples of annual data. It is also possible to allow for structural breaks in the context of a panel unit root process, as conducted by Wu *et al.* (2004), who reject the unit root null hypothesis for a set of Southeast Asian real exchange rates. However, the nature of our data including considerable differences in the lengths of the series and the sample periods covered mean that it is not possible to use a panel approach.

Perron (1997) proposes an extension of his original technique using a sequential procedure that estimates the test statistic allowing for a break at any point during the sample to be determined by the data. This technique is very similar to that of Zivot and Andrews, except that Perron's is more flexible, and therefore arguably preferable, since it allows for a break under both the null and alternative hypotheses whereas according to Zivot and Andrews' model, it can only arise under the alternative. Given its apparent superiority over the previous approaches, we employ the Perron (1997) test here, with the results being presented in Table 4.6.

[TABLE 4.6 HERE]

These are tests for whether the series are 'stationary' in their levels (rather than the percentage changes) – i.e. the question is whether they contain a stochastic trend, otherwise known as a unit root. The null hypothesis is that they do and asterisks denote instances where this hypothesis is rejected. The table presents results for both the pure unit root tests (i.e. the DF and ADF tests) and tests that allow for a single endogenously determined structural break (the Perron test).

As the results in the table show, there is a considerable degree of agreement between the test that includes an arbitrary three lags and the pure Dickey-Fuller (DF) test that does not. The evidence for modern foreign exchange series is more strongly that they are all non-stationary, whereas the results here are much more mixed. Almost half the series here appear to be stationary, including most of the Barcelona series and half of those from Venice.

The Perron test that we employ is for a unit root and a structural break, when the break date is assumed unknown. The null in each case is of a unit root with a break against an alternative of stationarity, and the results are very different to the DF tests. Allowing for a structural break, only two of the series are stationary in their levels, which is arguably much more the result that we would have expected. For the Venice-based series, there seems to be a number having breaks around September 1386 and September 1399 - April 1400. Several of the Barcelona series have breaks in August 1388, but here the patterns are not so clear. For rates quoted at Bruges, the most common date is between April and July 1396. This is interesting, in large part because it does not coincide with any of the frequent devaluations and *renforcements* of the Flemish currency (Munro, 2005). De Roover suggested that the rise in the Bruges-Venice rate may have been connected to the capture of John de Nevers, heir to the duchy of Burgundy (of which the county of Flanders was part) by the Ottomans at the battle of Nicopolis. The subsequent payment of John's ransom, set at 200,000 ducats, required large transfers from Bruges and Paris to Venice and thence to the east, which would have driven up the price of ducats (de Roover 1968, pp.52-3).

An important limitation of the Perron (1997) approach is that it can only be employed to estimate (up to) one structural break. This may be problematic not only in the sense that other breaks cannot be detected even if they are present in the data, but more seriously, if there are multiple breaks then a model assuming that there is at most one will be mis-specified, possibly leading to errors in inference as serious as those if there was one break and we allowed for none. Thus a further extension would be to allow for more than one structural break in the series – for example, Lumsdaine and Papell (1997) enhance the Zivot and Andrews (1992) approach to allow for two structural breaks. It is also possible to allow for structural breaks in the cointegrating relationship between series using an extension of the first step in the Engle-Granger approach – see Gregory and Hansen (1996).

More recently, in a series of papers, Bai and Perron (1998, 2003a, 2003b) develop another technique that tests for structural change, but this time allowing for more than one break. This approach, however, is not conducted in the context of unit root tests and is instead applied to a model using stationary data. In our case, in the absence of data on exogenous factors that may affect the exchange rates, the test for a structural break is based on a simple autoregressive model of order one (AR(1)) on the percentage changes in the exchange rates. Essentially, the test works by estimating a set of models allowing for 1, 2, ..., m structural breaks and selecting the number of breaks that minimises the Bayesian information criterion (BIC). It is also possible to use an F -test to examine the null hypothesis that there exists a given number of breaks m against the alternative of $m-1$. The model is thus

$$\begin{aligned}
 y_t &= x_t' \beta + z_t' \delta_1 + u_t, & t = 1, \dots, T_1 \\
 y_t &= x_t' \beta + z_t' \delta_2 + u_t, & t = T_1 + 1, \dots, T_2 \\
 &\vdots \\
 y_t &= x_t' \beta + z_t' \delta_{m+1} + u_t, & t = T_m + 1, \dots, T
 \end{aligned}
 \tag{4.11}$$

where y_t is the dependent variable, x_t is a $p \times 1$ vector of variables whose relationships with y are assumed not to vary over time, and z_t is a $q \times 1$ vector of variables whose relationships with y vary because of a set of structural breaks which take place at times T_1, T_2, \dots, T_m . The vectors $\beta, \delta_1, \dots, \delta_{m+1}$ represent sets of unknown parameter values. We conduct the test allowing for a maximum of five breaks, and a minimum distance between the breaks of three months.

The results from applying the test are given in Table 4.7. For around half of the series, no structural breaks are detected. Interestingly, almost all of the Barcelona series do not have any breaks, while all of the Bruges series except Bruges-Barcelona have a single break, and the results for the Venice pairs are more mixed and with occasional multiple breaks. Thus it appears that, overall, the possibility of multiple breaks need not be of significant concern here. Where breaks do take place, it is evident that there is much less consistency across series in their dates of occurrence than was the case for the unit root-structural break results presented above.

[TABLE 4.7 HERE]

One interesting case study is Venice-Milan. The Perron test reported in Table 4.6 identifies a structural break in April 1395 whereas the Bia-Perron break test reported in Table 4.7 found a structural break in March 1400. This coincides with two different periods of Milanese monetary policy (Mueller, 1997 pp.590-2). The exchange rate between Venice and Milan was quoted between the gold Venetian ducat and a notional gold Milanese ducat consisting of 32 silver *soldi imperiali*. Until 1395 the domestic exchange rate between ducats and *lire imperiali* in Milan remained stable and the exchange rate between Venice and Milan varied around 3-5 per cent in favour of Venice (i.e. 100 Venetian ducats were equivalent to 103-105 Milanese ducats/3296-3360 *soldi imperiali*). After 1395, however, Giangealeazzo Visconti, ruler of Milan, progressively debased the silver coinage to fund his aggressive foreign policy. As the value of the silver coin began to fall, so the Venice-Milan exchange rate rose. At its nadir in early 1400, the Milanese ducat was equivalent to 48-49 *soldi imperiali*, a decrease of roughly half in the value of the silver currency. In parallel, the Venice-Milan exchange rate increased to 52 per cent in favour of Venice (so that 100 Venetian ducats were now worth 152 Milanese ducats/4864 *soldi imperiali*). At this point, Visconti sought to revalue the silver currency in Milan and return to the 1395 valuation of 32 *soldi imperiali* to the Milanese ducat. On 21 February 1400, he issued a decree 'crying up' the value of the silver *lira imperiali* by one-third. By March 1400, the domestic exchange rate in Milan had fallen to 35-36 *soldi imperiali* per Milanese ducat and the Venice-Milan exchange rate had dropped correspondingly to 10-12 per cent in favour of Venice (i.e. 100 Venetian ducats were worth 110-112 Milanese ducats/3520-3584 *soldi imperiali*).

4.7. Conclusion

This chapter has introduced some fundamental statistical techniques that historians could use when analysing time series data. It has set out methods to identify seasonal trends, and to test for unit roots and structural breaks. Econometric analysis offers a number of advantages; it is more precise and, in some ways, more rigorous than merely 'eyeballing' the data. It can identify trends that a more superficial study of the raw numbers might miss or mis-interpret. It can determine whether, from a statistical perspective, an observed pattern is important or not. Moreover, a more formal quantitative analysis can also reduce the number of 'false positives' where the researcher effectively sees a face in the clouds that is not really there.

However, such statistical methods are not a replacement for history but need to be used in conjunction with traditional historical research. First, it is vital to consider the historical context when explaining any results produced by the statistical analysis. For example, structural break tests may propose potential dates but the ultimate meaning and significance of any such changes identified depends on the historical reconstruction. Second, and perhaps even more importantly, any statistical analysis is only as good as the data on which it is based. Historians must ensure the integrity of the data and that it is suitable for the analysis proposed.

References

Bai, Jushan and Perron, Pierre (1998) Estimating and testing linear models with multiple structural breaks, *Econometrica*, 66, 47-78

Bai, Jushan and Perron, Pierre (2003a) Computation and analysis of multiple structural change models, *Journal of Applied Econometrics*, 18, 1-22

Bai, Jushan and Perron, Pierre (2003b) Critical values for multiple structural change tests, *Econometrics Journal*, 1, 1-7

Banerjee, Anindya, Dolado, Juan J., Galbraith, John W. and Hendry, David F. (1993) *Cointegration, Error Correction, and the Econometric Analysis of Nonstationary Data* Oxford: Oxford University Press

Banerjee, Anindya, Lumsdaine, Ruben L. and Stock, James H. (1992) Recursive and sequential tests of the unit-root and trend-break hypotheses: Theory and international evidence, *Journal of Business and Economic Statistics*, 10, 271-287

Bell, Adrian R., Brooks, Chris, and Moore, Tony K. (2009) Interest in medieval accounts: Examples from England, 1272-1340, *History*, 94, 411-33

Bigwood, Georges and Grunzweig, Armand (eds) (1962) *Les Livres des Comptes des Gallerani* Brussels: Académie Royale de Belgique

Blomquist, Thomas W. (1990) Some observations on early foreign exchange banking based upon new evidence from XIIIth century Lucca, *Journal of European Economic History*, 19, 353-75

Booth, G. Geoffrey (2009) Foreign exchange profits in two early Renaissance money markets, *Journal of European Economic History*, 38, 123-144

Booth, G. Geoffrey and Gurun, Umit G. (2008) Volatility clustering and the bid-ask spread: Exchange rate behavior in early Renaissance Florence, *Journal of Empirical Finance*, 15, 133-144

Chilosi, David and Volckart, Oliver (2011) Money, states and empire: Financial integration and institutional change in Central Europe, 1400-1520, *Journal of Economic History*, 71, 762-91

- Christiano, Lawrence. J. (1992) Searching for a break in GNP, *Journal of Business and Economic Statistics*, 10, 237-250
- Corbae, Dean and Ouliaris, Sam (1988) Cointegration tests of purchasing power parity, *Review of Economics and Statistics*, 70, 508-511
- Denzel, Markus A. (2010) *Handbook of World Exchange Rates, 1590-1914*, Aldershot: Ashgate
- Dickey, David A. and Fuller, Wayne A. (1979) Distribution of estimators for time series regressions with a unit root, *Journal of the American Statistical Association*, 74, 427-431
- Einzig, Paul (1970) *The History of Foreign Exchange*, 2nd ed., London: Macmillan
- Fuller, W. A. (1976) *Introduction to Statistical Time Series*, New York: Wiley
- Goodhart, Charles A. E., McMahon, Patrick C. and Ngama, Yerima L. (1993) Testing for unit roots with very high frequency spot exchange rate data, *Journal of Macroeconomics* 15(3), 423-438
- Jorion, Philippe and Sweeney, Richard J. (1996) Mean reversion in real exchange rates: Evidence and implications for forecasting *Journal of International Money and Finance*, 15, 535-50
- Koyama, Mark (2010) Evading the 'Taint of Usury': the usury prohibition as a barrier to entry, *Explorations in Economic History*, 47, 420-42
- Kugler, Peter (2011) Financial market integration in late medieval Europe: Results from a threshold error correction model for the Rhinegulden and Basle Pound, 1365-1429, *Swiss Journal of Economics and Statistics*, 147, 337-57
- Leone, Alfonso (1983) Some preliminary remarks on the study of foreign currency exchange in the medieval period, *Journal of European Economic History*, 12, 619-29
- MacDonald, Ronald (1996) Panel unit root tests and real exchange rates, *Economics Letters*, 50, 7-11
- Meese, Richard A. and Singleton, Kenneth J. (1982) On unit roots and the empirical modelling of exchange rates *Journal of Finance*, 37(4), 1029-1035
- Melis, Federigo (1962) *Aspetti della Vita Economica Medievale: Studi nell'Archivio Datini di Prato*, Siena: Monte dei Paschi di Siena
- McCusker, John J. and Gravesteyn, Cora (1991) *The Beginnings of Commercial and Financial Journalism: The Commodity Price Currents, Exchange Rate Currents, and Money Currents of Early Modern Europe*, Amsterdam: Nederlandsch Economisch-Historisch Archief
- Mueller, Reinhold C. (1997) *The Venetian Money Market: Banks, Panics, and the Public Debt, 1200-1500*, Baltimore, MD: Johns Hopkins Press

Munro, John H. A. (1979) Bullionism and the bill of exchange in England, 1272-1663 in *The Dawn of Modern Banking: Ten Essays Prepared for a Conference Sponsored by the Centre for Medieval and Renaissance Studies, University of California, Los Angeles* London: Yale University Press, pp.169-239.

Munro, John H. A. (2005) Coinage debasements in Burgundian Flanders, 1384-1482: Monetary or fiscal remedies? in Nicholas, David, Bachrach, Bernard S., and Murray, James M. (eds.) *Comparative Perspectives on History and Historians: Essays in Honour of Bryce Lyon*, Kalamazoo: Medieval Institute Publications, Western Michigan University, pp. 314-60.

Nigro, Giampiero (ed.) (2010) *Francesco di Marco Datini: The Man, the Merchant*, Florence: Firenze University Press, Fondazione Istituto Internazionale di Storia Economica 'F. Datini'

Oh, Keun-Yeob (1996). Purchasing power parity and unit root tests using panel data, *Journal of International Money and Finance*, 15, 405-418

Origo, Iris M. (1963) *The Merchant of Prato: Francesco di Marco Datini, 1335-1410* rev. ed. Harmondsworth: Penguin

Pagnini, Gianfrancesco (ed.) (1766) *Della Decima e di Varie Altre Gravezze Imposte dal Comune di Firenze della Monete e della Mercatura de Fiorentini Fino al Secolo XVI: Tomo Quarto Continente La Pratica della Mercatura Scritta da Giovanni di Antonio da Uzzano nel 1442*, Lisbon-Lucca

Perron, Pierre (1997) Further evidence on breaking trend functions in macroeconomic variables *Journal of Econometrics*, 80, 355-385

Perron, Pierre (1989) The Great Crash, the Oil Price Shock and the Unit Root Hypothesis *Econometrica*, 57, 1361-1401

Quandt, Richard (1960) Tests of the hypothesis that a linear regression system obeys two different regimes, *Journal of the American Statistical Association*, 55, 324-330

de Roover, Raymond A. (1944a) Early accounting problems of foreign exchange, *The Accounting Review*, 19, 381-407

de Roover, Raymond A. (1944b) What is dry exchange? A contribution to the study of English Mercantilism, *Journal of Political Economy*, 52, 250-66

de Roover, Raymond A. (1968) *The Bruges Money Market Around 1400 (with a Statistical Supplement by Hyman Sardy)*, Brussels: Paleis der Academiën

Rubin, Jared (2010) Bills of exchange, interest bans and impersonal exchange in Islam and Christianity, *Explorations in Economic History*, 47, 213-27.

Schotman, Peter and Dijk, Herman K. van (1991) A Bayesian analysis of the unit root in real exchange rates, *Journal of Econometrics*, 49, 195-238

Spufford, Peter (1986) *Handbook of Medieval Exchange*, London: Royal Historical Society

Spufford, Peter (1991) *Money and its Use in Medieval Europe*, Cambridge: Cambridge University Press

Taylor, Mark P. (1990) On unit roots and real exchange rates: Empirical evidence and Monte Carlo analysis, *Applied Economics*, 22, 1311-1321

Volckart, Oliver and Wolf, Nikolaus (2006) Estimating financial integration in the Middle Ages: What can we learn from a TAR Model? *Journal of Economic History*, 66, 122-39

Whitt, Joseph A. Jr. (1992) The long-run behaviour of the real exchange rate: A reconsideration, *Journal of Money, Credit and Banking*, 24(1), 72-82

Wu, Yangru (1996) Are real exchange rates nonstationary? Evidence from a panel data test, *Journal of Money, Credit and Banking*, 28, 54-63

Wu, Jyh-Lin, Tsai, Li-Ju and Chen, Show-Lin (2004) Are real exchange rates non-Stationary? The Pacific Basin perspective, *Journal of Asian Economics*, 15, 425-438

Zivot, Eric and Andrews, Donald W.K. (1992) Further evidence on the Great Crash, the Oil Price Shock, and the Unit Root Hypothesis, *Journal of Business and Economic Statistics*, 10, 251-270

Table 4.1: Definitions of currency pairs and available sample

Series	Where Quoted	Meaning	Total Observations	Main Date Range	No. of interpolated values
Barcelona-Avignon	Barcelona	Pence of Barcelona per French franc	229	1386:11 – 1405:11	25
Barcelona-Bruges	Barcelona	Pence of Barcelona per écu of 22 groats	224	1387:01 – 1405:11	20
Barcelona-Florence	Barcelona	Pence of Barcelona per Florentine florin	102	1394:10 – 1405:11	6
Barcelona-Genoa	Barcelona	Pence of Barcelona per florin of Genoa	229	1386:11 – 1405:11	19
Barcelona-Majorca	Barcelona	Pence of Barcelona per Majorcan real	132	1390:02 – 1405:09	19
Barcelona-Montpellier	Barcelona	Pence of Barcelona per French franc	218	1387:01 – 1405:11	15
Barcelona-Pisa	Barcelona	Pence of Barcelona per French franc	192	1387:01 – 1405:08	6
Barcelona-Venice	Barcelona	Pence of Barcelona per Venetian ducat	80	1399:02 – 1405:11	1
Bruges-Barcelona	Bruges	Pence of Barcelona per écu of 22 groats	194	1392:01 – 1410:09	5
Bruges-Genoa	Bruges	Flemish groats per florin of Genoa	196	1389:03 – 1410:09	12
Bruges-London	Bruges	Pence sterling per écu of 24 groats	190	1389:03 – 1410:09	12
Bruges-Paris	Bruges	Flemish groats per French franc	195	1389:03 – 1410:08	10
Bruges-Pisa	Bruges	Flemish groats per florin of Pisa	39	1395:03 – 1397:08	3
Bruges-Venice	Bruges	Flemish groats per Venetian ducat	196	1389:03 – 1410:09	14
Venice-Barcelona	Venice	Pence of Barcelona per Venetian ducat	139	1399:03 – 1410:09	1
Venice-Bologna	Venice	Florins of Bologna per Venetian ducat	269	1384:05 – 1410:09	14
Venice-Bruges	Venice	Flemish groats per Venetian ducat	251	1384:03 – 1410:09	2
Venice-Florence	Venice	Pence affiorino of Florence per Venetian ducat	320	1384:02 – 1410:09	10
Venice-Genoa	Venice	Florins of Genoa per Venetian ducat	319	1384:03 – 1410:09	16
Venice-London	Venice	Pence sterling per Venetian ducat	65	1403:10 – 1410:09	7
Venice-Lucca	Venice	Florins of Lucca per Venetian ducat	135	1399:04 – 1410:06	10
Venice-Milan	Venice	Ducats of Milan per Venetian ducat	136	1384:03 – 1404:03	6
Venice-Paris	Venice	Grossi d'oro of Venice per French franc	177	1388:08 – 1410:09	7
Venice-Pisa	Venice	Florins of Pisa per Venetian ducat	266	1384:03 – 1410:05	14
Venice-Rome	Venice	Florins of Rome per Venetian ducat	99	1394:06 – 1407:08	1

Table 4.2: Summary statistics for exchange rate returns (percentage changes)

Series	Mean	Variance	Skewness	Kurtosis (excess)	Min	Max
Barcelona-Avignon	-0.01	1.96	0.20	1.54	-4.45	4.94
Barcelona-Bruges	0.09	13.02	9.35	119.04	-6.45	46.15
Barcelona-Florence	0.00	2.60	0.54	2.82	-3.92	6.80
Barcelona-Genoa	0.01	3.48	-0.53	2.12	-8.42	5.94
Barcelona-Majorca	0.05	0.78	-0.45	3.52	-3.95	2.84
Barcelona-Montpellier	0.00	2.18	0.33	2.41	-4.57	6.11
Barcelona-Pisa	0.02	3.39	-0.10	3.10	-7.03	8.50
Barcelona-Venice	0.02	1.90	-0.12	0.29	-3.24	3.94
Bruges-Barcelona	-0.05	2.28	0.01	1.35	-4.99	4.53
Bruges-Genoa	0.00	2.50	-0.28	1.00	-5.64	4.49
Bruges-London	-0.04	1.43	0.44	1.54	-4.08	3.85
Bruges-Paris	0.06	0.81	-0.23	2.42	-3.4	3.41
Bruges-Pisa	0.12	1.47	-0.44	1.20	-3.47	2.50
Bruges-Venice	0.12	9.49	2.68	62.85	23.08	30.3
Venice-Barcelona	0.02	1.99	-0.11	1.50	-5.00	4.35
Venice-Bologna	0.00	0.57	0.02	3.63	-2.97	3.51
Venice-Bruges	0.04	2.22	-0.23	2.22	-6.77	4.83
Venice-Florence	0.01	1.07	-0.08	2.78	-3.96	5.19
Venice-Genoa	0.07	1.48	-0.08	1.34	-3.96	3.98
Venice-London	0.14	1.76	0.48	1.50	-2.67	4.44
Venice-Lucca	0.01	0.94	0.10	4.98	-4.32	3.93
Venice-Milan	0.10	5.97	-5.27	48.98	21.85	7.42
Venice-Paris	-0.03	1.86	-0.13	0.60	-4.08	4.34
Venice-Pisa	0.02	1.10	0.29	4.09	-3.88	5.40
Venice-Rome	0.20	2.26	1.03	3.81	-3.04	7.03

Table 4.3: Seasonal variations – Barcelona

	Avignon	Bruges	Florence	Genoa	Majorca	Montpellier	Paris	Venice
January	0.87*** (0.30)	1.68** (0.81)	1.63*** (0.53)	1.85*** (0.39)	0.84*** (0.24)	0.94*** (0.33)	2.17*** (0.42)	0.85 (0.53)
February	0.39 (0.30)	1.06 (0.81)	0.75 (0.53)	0.86** (0.39)	0.61*** (0.23)	0.52 (0.32)	0.50 (0.42)	0.81 (0.49)
March	-0.02 (0.30)	0.16 (0.83)	-0.35 (0.53)	0.14 (0.39)	0.12 (0.25)	-0.02 (0.33)	0.54 (0.44)	0.39 (0.49)
April	0.22 (0.30)	0.29 (0.83)	0.59 (0.53)	0.31 (0.39)	0.22 (0.24)	0.22 (0.33)	0.38 (0.42)	0.86* (0.49)
May	0.23 (0.30)	0.04 (0.83)	-1.16** (0.53)	-0.90** (0.39)	0.45* (0.25)	0.24 (0.32)	-0.29 (0.41)	-0.91* (0.49)
June	0.34 (0.30)	-0.21 (0.81)	-0.55 (0.53)	-0.71* (0.39)	-0.02 (0.24)	0.26 (0.32)	-0.47 (0.41)	-0.23 (0.49)
July	-0.43 (0.30)	-0.77 (0.81)	-0.13 (0.50)	-0.68* (0.39)	-0.25 (0.24)	-0.24 (0.33)	-0.80* (0.41)	-0.54 (0.49)
August	-0.53* (0.30)	-0.94 (0.81)	-0.48 (0.50)	-0.47 (0.39)	-0.14 (0.24)	-0.84** (0.33)	-0.36 (0.42)	-0.75 (0.49)
September	-1.17*** (0.30)	-1.8** (0.81)	-1.20** (0.50)	-0.72* (0.39)	-0.16 (0.23)	-1.18*** (0.33)	-1.04* (0.44)	-1.12** (0.53)
October	-0.84*** (0.30)	-1.1 (0.81)	0.20 (0.47)	-0.65* (0.39)	-0.81*** (0.22)	-0.67** (0.33)	-0.65 (0.42)	0.1 (0.53)
November	0.07 (0.29)	2.17*** (0.81)	0.04 (0.50)	-0.15 (0.38)	-0.31 (0.25)	0.08 (0.33)	-0.09 (0.42)	0.42 (0.49)
December	0.71** (0.30)	0.54 (0.83)	0.77 (0.53)	1.28*** (0.39)	0.19 (0.25)	0.67** (0.34)	0.46 (0.44)	0.38 (0.53)
R ²	0.18	0.10	0.24	0.21	0.25	0.17	0.20	0.24

Notes: Standard errors in parentheses; *, ** and *** denote significance at the 10 per cent, 5 per cent, and 1 per cent levels; critical values vary according to the number of observations for each series. All the rates are quoted as an uncertain number of pence of Barcelona per unit of foreign currency.

Table 4.4: Seasonal variations – Bruges

	Barcelona	Genoa	London	Paris	Pisa	Venice
January	-0.23 (0.37)	1.09*** (0.36)	-0.57** (0.28)	0.39* (0.22)	1.16* (0.66)	0.82 (0.76)
February	-0.65* (0.36)	0.56 (0.35)	-0.58** (0.28)	0.11 (0.21)	0.25 (0.57)	0.49 (0.73)
March	-0.58 (0.36)	0.33 (0.34)	-0.49* (0.28)	0.42** (0.21)	0.51 (0.51)	0.60 (0.71)
April	0.41 (0.36)	-0.92*** (0.34)	-0.53** (0.27)	-0.25 (0.21)	-0.15 (0.51)	-0.82 (0.71)
May	0.3 (0.36)	-0.51 (0.35)	-0.55* (0.28)	0.01 (0.21)	-0.02 (0.57)	-0.25 (0.73)
June	0.08 (0.36)	-1.38*** (0.35)	0.14 (0.28)	-0.04 (0.21)	-1.23** (0.57)	-0.40 (0.73)
July	0.11 (0.37)	-0.48 (0.36)	0.12 (0.28)	-0.58*** (0.22)	-0.78 (0.66)	-0.68 (0.76)
August	-0.4 (0.37)	0.41 (0.36)	0.03 (0.28)	0.12 (0.22)	0.38 (0.66)	-1.52** (0.76)
September	-0.56 (0.37)	0.08 (0.36)	0.96*** (0.28)	0.36 (0.23)	1.15 (0.81)	2.12*** (0.76)
October	0.21 (0.38)	0.40 (0.38)	0.84*** (0.29)	0.17 (0.23)	1.49* (0.81)	0.21 (0.78)
November	0.95** (0.38)	-0.12 (0.38)	0.37 (0.29)	-0.14 (0.23)	-0.16 (0.81)	-0.01 (0.78)
December	-0.11 (0.38)	0.80** (0.38)	-0.15 (0.29)	0.18 (0.23)	-0.15 (0.81)	1.09 (0.78)
R ²	0.09	0.20	0.19	0.10	0.37	0.09

Notes: The rates for Genoa, Paris, Pisa and Venice are quoted as an uncertain number of Flemish groats per unit of foreign currency while Barcelona and London are quoted as an uncertain number of foreign coins per écu (of 22 and 24 groats respectively).

Table 4.5: Seasonal variations – Venice

	Barcelona	Bologna	Bruges	Florence	Genoa	London	Lucca	Milan	Paris	Pisa	Rome
January	0.05 (0.40)	-0.15 (0.15)	0.11 (0.35)	-0.19 (0.18)	-0.33 (0.22)	-0.64 (0.59)	0.14 (0.27)	0.74 (0.68)	0.38 (0.32)	-0.19 (0.20)	-0.01 (0.57)
February	-0.46 (0.40)	-0.09 (0.15)	-1.03*** (0.36)	-0.45** (0.18)	-0.33 (0.22)	-0.51 (0.59)	-0.40 (0.27)	-0.70 (0.68)	0.80** (0.33)	-0.46** (0.20)	-0.05 (0.57)
March	-0.68* (0.38)	0.04 (0.15)	0.11 (0.34)	0.18 (0.18)	0.10 (0.22)	-0.24 (0.59)	0.12 (0.27)	-0.66 (0.68)	0.20 (0.33)	0.10 (0.19)	-0.13 (0.53)
April	-0.29 (0.38)	0.01 (0.15)	0.04 (0.34)	-0.04 (0.18)	0.20 (0.22)	-0.45 (0.54)	0.29 (0.26)	0.40 (0.78)	-0.32 (0.34)	0.01 (0.19)	0.09 (0.57)
May	0.01 (0.38)	-0.18 (0.15)	0.01 (0.35)	0.13 (0.18)	0.64*** (0.22)	1.14** (0.54)	0.09 (0.26)	0.16 (0.78)	0.04 (0.34)	0.01 (0.19)	0.16 (0.50)
June	0.86** (0.38)	0.02 (0.15)	1.37*** (0.33)	0.56*** (0.18)	0.62*** (0.22)	0.57 (0.54)	0.60** (0.26)	1.02 (0.74)	-1.01*** (0.33)	0.71*** (0.20)	0.29 (0.45)
July	0.77** (0.38)	0.39** (0.15)	0.60* (0.34)	0.58*** (0.18)	0.57*** (0.22)	0.55 (0.54)	0.67** (0.27)	0.46 (0.74)	-0.89*** (0.33)	0.91*** (0.20)	1.11** (0.47)
August	0.33 (0.38)	0.20 (0.15)	-0.04 (0.33)	-0.24 (0.18)	0.17 (0.22)	0.56 (0.54)	-0.71*** (0.27)	-0.10 (0.74)	-0.46 (0.32)	-0.16 (0.19)	0.81* (0.45)
September	-1.27*** (0.38)	-0.64*** (0.15)	-0.81** (0.33)	-1.29*** (0.18)	-1.11*** (0.22)	-0.30 (0.59)	-0.83*** (0.27)	-1.02 (0.74)	0.81** (0.33)	-1.35*** (0.20)	-0.76 (0.53)
October	0.42 (0.40)	0.01 (0.15)	0.66* (0.34)	0.51*** (0.18)	0.39* (0.22)	0.24 (0.59)	0.28 (0.27)	0.38 (0.78)	-0.39 (0.34)	0.39* (0.20)	0.61 (0.57)
November	-0.28 (0.40)	0.17 (0.15)	0.07 (0.34)	0.32* (0.18)	-0.03 (0.22)	0.68 (0.59)	-0.20 (0.27)	-0.04 (0.78)	0.57* (0.34)	0.16 (0.20)	-0.43 (0.57)
December	0.78* (0.40)	0.17 (0.15)	-0.10 (0.34)	0.08 (0.18)	-0.14 (0.22)	-0.22 (0.59)	-0.01 (0.27)	0.69 (0.68)	-0.09 (0.34)	0.02 (0.20)	0.07 (0.57)
R ²	0.20	0.11	0.14	0.24	0.16	0.18	0.21	0.07	0.19	0.27	0.13

Notes: All the exchange rates are quoted as an uncertain number of foreign coins per Venetian ducat, except for Venice-Paris.

Table 4.6: Unit Root and Structural Break Tests

Series	DF t-statistic (0 lags)	ADF t-statistic (3 lags)	Perron break test t-statistic	Break date
Barcelona-Avignon	-4.69***	-4.81***	-0.46	1388:08
Barcelona-Bruges	-2.06	-2.40	-0.82	1389:02
Barcelona-Florence	-2.31	-2.46	-0.45	1398:04
Barcelona-Genoa	-3.95***	-5.01***	0.08	1388:08
Barcelona-Majorca	-4.25***	-5.68***	-1.43	1393:01
Barcelona-Montpellier	-4.81***	-4.85***	-1.29	1389:10
Barcelona-Pisa	-2.78*	-3.11**	-1.20	1393:07
Barcelona-Venice	-1.55	-2.21	-0.86	1399:08
Bruges-Barcelona	-3.10***	-2.39	-0.60	1396:07
Bruges-Genoa	-3.12**	-3.09**	-0.34	1396:07
Bruges-London	-2.33	-2.29	-3.95***	1396:09
Bruges-Paris	-2.23	-1.33	0.62	1396:07
Bruges-Pisa	-0.98	-1.27	-0.83	1395:04
Bruges-Venice	-4.50***	-2.71*	-0.45	1396:07
Venice-Barcelona	-1.87	-1.64	-0.60	1400:03
Venice-Bologna	-6.12***	-3.85***	0.11	1392:11
Venice-Bruges	-1.49	-1.30	-0.82	1399:10
Venice-Florence	-6.05***	-3.98***	0.01	1386:09
Venice-Genoa	-1.18	-0.78	1.04	1386:09
Venice-London	-0.77	-0.48	-3.10**	1406:08
Venice-Lucca	-4.17***	-3.02**	-1.06	1400:04
Venice-Milan	-2.4	-2.12	-0.88	1395:04
Venice-Paris	-1.46	-1.44	-1.25	1400:02
Venice-Pisa	-6.30***	-5.34**	-1.16	1386:03
Venice-Rome	-2.81*	-2.60	-1.17	1399:09

Notes: A constant but no trend is included in all DF and ADF test regressions; the additive outlier model is used for the Perron test.

Table 4.7: Bai-Perron structural break tests

Series	Number of Structural Breaks	Break date	Second break date	Third break date
Barcelona-Avignon	0			
Barcelona-Bruges	0			
Barcelona-Florence	0			
Barcelona-Genoa	0			
Barcelona-Majorca	0			
Barcelona-Montpellier	1	1395:06		
Barcelona-Pisa	2	1399:02	1398:12	
Barcelona-Venice	0			
Bruges-Barcelona	0			
Bruges-Genoa	1	1396:07		
Bruges-London	1	1407:12		
Bruges-Paris	1	1395:11		
Bruges-Pisa	1	1396:07		
Bruges-Venice	1	1410:03		
Venice-Barcelona	0			
Venice-Bruges	0			
Venice-Bologna	1	1403:07	1400:03	
Venice-Florence	0			
Venice-Genoa	0			
Venice-London	1	1406:11		
Venice-Lucca	0	1404:07	1409:12	
Venice-Milan	2	1400:03	1399:09	
Venice-Paris	1	1409:03		
Venice-Pisa	3	1397:10	1397:06	1398:01
Venice-Rome	1	1403:10		

Note: The number of structural breaks is selected by the Bai-Perron test that minimises the BIC; the tests are constructed in the context of an AR(1) model.