The JPY/AUD Carry Trade and Its Causal Linkages to Other Markets

Brian D. Deaton
McMurry University

This study analyzes the causal structure underlying the popular Japanese Yen/Australian Dollar (JPY/AUD) carry trade and related financial variables. Three causal search algorithms are employed to find the relationships amongst the JPY/AUD exchange rate, the S&P 500 stock index, the Nikkei 225 stock index, the Australian Securities Exchange 200 stock index, the 10-year U.S. Treasury Note, the 10-year Japanese government bond, and the 10-year Australian government bond. The results from all three algorithms provide evidence against the theory of uncovered interest rate parity.

Keywords: currency carry trade, uncovered interest rate parity, causality, vector autoregression, market linkages

INTRODUCTION

Long used by hedge funds and institutional investors, currency carry trade investment strategies are now becoming popular with individual investors. Several exchange-traded vehicles, such as the Invesco DB G10 Harvest Fund and the iPath Optimized Currency Carry ETN, have been developed to make it easy for the small investor to participate in carry trade strategies. The popularity of carry trade strategies coupled with ever increasing financial integration between financial markets might lead to spillover effects from currencies to stocks and bonds or vice versa. It is the goal of this paper look for these types of linkages between markets.

A currency carry trade is constructed by borrowing money in a currency with low interest rates (funding currency) and simultaneously investing that money in a currency with higher interest rates (target currency) with the goal of profiting on the interest rate differential. This strategy is not without risk because it is possible that adverse exchange rate movements could quickly erase the investor's profits. In fact, it has been said that "Trading currency carry trades is like picking up nickels in front of steamrollers: you have a long run of small gains but eventually get squashed" ("Carry on Speculating," 2007).

Traditionally, the Japanese yen has been an extremely popular funding currency because of its very low interest rate while the Australian dollar has been a popular target currency because of its relatively high interest rate (Fung, Tse, & Zhao, 2013). Therefore, this study focuses on the Japanese yen to Australian dollar (JPY/AUD) currency pair and attempts to find linkages between this exchange rate and the stock and bond markets of Japan, Australia, and the United States. Market linkages may exist due to market participants seeking portfolio diversification or perhaps due to some other factor such as an information flow contained within price movements.

Economic Theory

Theory asserts that currency carry trades should not be profitable. According to uncovered interest rate parity, the exchange rate should change just enough to offset any interest rate differential. This is summarized by

$$i_H - i_L = \Delta s_{t+1} \tag{1}$$

where $\boldsymbol{i_H}$ is the interest rate of the high-yielding currency, $\boldsymbol{i_L}$ is the interest rate of the low yielding currency, and $\Delta \boldsymbol{s_{t+1}}$ is the change in the spot exchange rate in units of high-yielding currency per unit of low-yielding currency (H/L) within the next time period ($\boldsymbol{t}$ is a time index) (Bekaert & Hodrick, 2012). Although uncovered interest rate parity is in virtually every international economics textbook, it seldom actually holds. Some studies estimate that the currency carry trades are profitable up to eighty percent of the time.

LITERATURE REVIEW

Other studies have found evidence that currency carry trade returns have a causal relationship with other financial markets such as stock or bond markets. A brief summary of the literature is provided below.

Using daily data, (Fung et al., 2013) identifies significant Granger causal relationships from carry trade returns to the stock markets of Japan, Australia, and India during the 2008 financial crisis using constructed carry trade baskets of G10 currencies (i.e. the ten most liquid currencies). Interestingly, the study finds that no causal relationship exists in the pre-crisis period from January 1995 to July 2007.

(Cheung, Cheung, & He, 2012) uses weekly data and three different proxies for currency carry trade activity to conclude that the yen carry trade affects stock markets in five chosen target currency countries while controlling for US stock returns, the VIX index, and commodity prices.

(Tse & Zhao, 2012) uses daily data from 1995-2010 with both a VAR model and an EGARCH-t model to determine that the U.S. stock market and carry trade returns are correlated and that there is volatility spillover from the stock market to the currency market. The study searches for Granger-causality but finds none in either direction.

(Zhang, Yau, & Fung, 2010) uses a VAR models with daily data for several economic variables including four currency markets, stock markets, interest rates, and credit default swaps. The study finds significant Granger-causality flowing from the credit default swap market to the JPY/USD, AUD/USD, and EUR/USD. The conclusion is that the credit market influences the currency market via carry trades but not vice-versa.

(Christiansen, Ranaldo, & Söderlind, 2010) uses daily data from 1995-2008 to investigate the exposure of G10 currency carry trade returns to the S&P 500 futures and the 10-year U.S. T-note futures contracts during different volatility regimes. The study finds that carry trade returns are significantly positively exposed to stock returns during all periods, but during periods of turmoil, the exposure is much larger. The exposure of carry trade returns to the bond market is found to be negligible and not significant.

Using five-minute intraday data from 1993-2008, (Ranaldo & Söderlind, 2010) find safe haven properties of the Swiss franc, Japanese yen, and, to a lesser extent, the euro in that their excess returns are negatively related to the S&P 500 stock index but positively related to the 10-Year US Treasury note and foreign exchange volatility.

(Nishigaki, 2007) looks at five financial variables including the U.S. stock market and the Japanese stock market using monthly data to investigate the yen carry trade. The study finds that a change in the U.S. stock market affects yen futures contract positions on the Chicago Mercantile Exchange.

In summary, six of the cited studies find a connection between the carry trade and one or more stock markets. Only one study, (Ranaldo & Söderlind, 2010), finds a connection between the carry trade and a bond market.

METHODOLOGY

Causality

The tried and true method for discovering causal relationships is the controlled experiment. In many cases, such as in this study where we try to discover the causal linkages between financial markets, a controlled experiment cannot be performed. Fortunately, methods of causal discovery based on observational data have recently been developed.

Causal models are insightful because they contain more information than the more common probabilistic models in that they represent the data generating process and not just the statistical regularities of the data. While probabilistic models can be used to predict one variable from the observation of another, causal models can, in addition, allow the prediction of one variable when intervening on another. Determining the effect of an intervention is especially helpful for policy makers when trying to predict the potential effects of a policy. Practitioners in the financial markets may also find this useful when trying to determine the price effects of some intervention by a large institution or government body, e.g. a central bank.

The meaning of causation in this case is simply when the value of one variable is changed and that of some other variable also changes; the former is the cause and the latter the effect. Causal relationships may be encoded in directed acyclic graphs (DAGs), which present a nice visual representation of the causal connections.

Directed Acyclic Graphs

A directed graph is a diagram that shows a set of vertices connected by a set of directed edges. For example, for a pair of variables A and B that are represented in the graph by a pair of vertices with the same names, if a directed edge connects A to B, shown as $A \rightarrow B$, then A has a causal influence on B. In a directed acyclic graph, no path (sequence of directed edges) can lead from a vertex back to itself. That is, a DAG contains no directed cycles, e.g. $A \rightarrow B \rightarrow A$, or A cannot be a cause of itself.

Sometimes one DAG does not uniquely identify the causal structure underlying a set of variables. The conditional independence information contained in a joint probability distribution could possibly be represented by more than one DAG. In this case, a set of DAG structures that are observationally equivalent is presented. An equivalence class of DAGs is often represented by a partially directed DAG in which some edges are not directed. The directed edges are common to all members of the equivalence class while the undirected edges point one way in some members of the equivalence class and the opposite way in other members. See (Pearl, 2000) for more information on causality and DAGs.

Causal Search Algorithms

To find the connections between currency, bond, and stock markets, this study uses three causal search algorithms. The first two algorithms, Linear Non-Gaussian Acyclic Model (LiNGAM) and Fast Causal Search (FCI), are not meant for time series data and must be used in conjunction with a vector autoregression (VAR) model. The third algorithm, FCI for time series (tsFCI) uses time series data directly. All three algorithms are included in the Tetrad software package (Glymour, Scheines, Spirtes, & Ramsey, 2017).

The LiNGAM algorithm assumes that observed variables are linear functions of disturbance variables and that the disturbance terms are mutually independent with non-Gaussian distributions and non-zero variances. In addition, it assumes that there are no latent common causes. The non-Gaussian assumption is key and allows LiNGAM to use independent component analysis (ICA) as a first phase. After performing independent component analysis, LiNGAM permutes and normalizes the ICA mixing matrix to obtain the causal ordering of the observed variables and the strength of their connections. This information can easily be transformed into a DAG to display the causal structure visually. For further information, see (Shimizu, Hoyer, Hyvärinen, & Kerminen, 2006).

The FCI algorithm allows for the possible existence of latent variables that affect the observed variables. FCI's output is a partial ancestral graph (PAG), which is slightly different from a DAG in that it represents the ancestor relationships shared by an equivalent class of DAGs. In a PAG,

1. A → B only if A is an ancestor of B
2. A → B only if B is not an ancestor of A (∘ is any type of endpoint)
3. A ↔ B implies that A and B may share a latent common cause

For a further explanation of the FCI algorithm, see (Spirtes, Glymour, & Scheines, 2000).

FCI for time series (tsFCI) is a modification of the FCI algorithm that incorporates knowledge that the data came from a time-series process. In particular, it acknowledges that causal effects must go forward in time. It also assumes the data is generated by a first-order time-invariant multivariate Markov process with sparse connections and possible hidden variables. See (Entner & Hoyer, 2010) for further information.

Description of the Data

Stock index data for the Australian Securities Exchange 200 index (AUS200), Nikkei 225 index (JPN225), and S&P 500 index (SPX500) was obtained from the Sierra Chart data service (*Sierra Chart*, 2019). Data for the JPY/AUD spot foreign exchange rate (JPY/AUD), Australian 10-year treasury bond futures contract (XS), Japanese 10-year government bond futures contract (JGB), and U.S. 10-year treasury-note futures contract (ZN) was obtained from the CQG data service (*CQG*, 2019). The data spans the timeframe 2017/11/1 – 2018/10/30 and has 10 minute periodicity. Sierra Chart software was used to join each bond's future contracts into a single continuous time series. Log returns, computed by taking the natural logarithm and first-differencing (in that order), of each variable are used in the analysis that follows.

RESULTS

VAR Estimation

Stata (*Stata*, 2019) was used to select and estimate an order one vector autoregression containing all of the observed variables (AUS200, JPN225, SPX500, JPY/AUD, XS, JGB, and ZN). The Hannan-Quinn information criterion and the Schwarz's Bayesian criterion both indicated that an order one (lag length one) VAR model is appropriate. Table 1 contains the VAR estimation result.

TABLE 1
VECTOR AUTOREGRESSION ESTIMATION RESULTS

JPY/AUD		*Coef.*	*Std. Err.*	*z*	*P>|z|*	*[95% Conf. Interval]*	
	JPY/AUD:1	0.996916	0.000475	2097.81	0	0.995985	0.997848
	AUS200:1	1.37E-06	3.36E-06	0.41	0.683	5.22E-06	7.96E-06
	XS:1	-5.69E-06	9.30E-07	-6.11	0	7.51E-06	-3.86E-06
	JGB:1	-2.06E-06	1.83E-05	-0.11	0.91	3.79E-05	3.38E-05
	JPN225:1	-2.76E-07	8.06E-07	-0.34	0.732	1.85E-06	1.30E-06
	SPX500:1	1.23E-05	8.67E-06	1.42	0.156	4.68E-06	2.93E-05
	ZN:1	0.005391	0.000859	6.27	0	0.003707	0.007076
	constant	5.137666	0.847634	6.06	0	3.476334	6.798998

AUS200	*Coef.*	*Std. Err.*	*z*	*P>|z|*	*[95% Conf. Interval]*	
JPY/AUD:1	-0.11382	0.038641	-2.95	0.003	0.189558	-0.03809
AUS200:1	1.000008	0.000273	3658.1	0	0.999472	1.000544
XS:1	-0.00025	7.56E-05	-3.32	0.001	0.000399	-0.0001
JGB:1	0.002987	0.001488	2.01	0.045	7.13E-05	0.005902
JPN225:1	4.22E-05	6.55E-05	0.64	0.519	8.61E-05	0.000171
SPX500:1	0.001216	0.000705	1.72	0.085	0.000166	0.002597
ZN:1	0.240714	0.069883	3.44	0.001	0.103746	0.377681
constant	175.4468	68.92223	2.55	0.011	40.36175	310.5319

XS	*Coef.*	*Std. Err.*	*z*	*P>|z|*	*[95% Conf. Interval]*	
JPY/AUD:1	-1.80034	0.482648	-3.73	0	2.746317	-0.85437
AUS200:1	0.001485	0.003415	0.43	0.664	0.005207	0.008178
XS:1	0.995273	0.000945	1053.59	0	0.993422	0.997125
JGB:1	-0.01389	0.01858	-0.75	0.455	0.050308	0.022525
JPN225:1	-0.00077	0.000818	-0.94	0.345	0.002376	0.000831
SPX500:1	0.008028	0.008803	0.91	0.362	0.009226	0.025282
ZN:1	3.424134	0.872886	3.92	0	1.71331	5.134958
constant	4532.878	860.8893	5.27	0	2845.566	6220.19

JGB	*Coef.*	*Std. Err.*	*z*	*P>|z|*	*[95% Conf. Interval]*	
JPY/AUD:1	0.011328	0.011424	0.99	0.321	0.011064	0.033719
AUS200:1	9.52E-05	8.08E-05	1.18	0.239	6.32E-05	0.000254
XS:1	-1.89E-06	2.24E-05	-0.08	0.933	4.57E-05	4.19E-05
JGB:1	0.998015	0.00044	2269.28	0	0.997153	0.998877
JPN225:1	-3.6E-05	1.94E-05	-1.88	0.06	7.43E-05	1.57E-06
SPX500:1	-0.00064	0.000208	-3.06	0.002	0.001047	-0.00023
ZN:1	-0.01729	0.020661	-0.84	0.403	0.057785	0.023207
constant	34.88612	20.37742	1.71	0.087	5.052881	74.82512

JPN225	*Coef.*	*Std. Err.*	*z*	*P>|z|*	*[95% Conf. Interval]*	
JPY/AUD:1	-0.71399	0.224231	-3.18	0.001	1.153476	-0.27451
AUS200:1	0.000352	0.001586	0.22	0.825	0.002758	0.003461
XS:1	-0.00173	0.000439	-3.94	0	0.002587	-0.00087
JGB:1	0.005628	0.008632	0.65	0.514	0.01129	0.022547
JPN225:1	0.999243	0.00038	2628.81	0	0.998498	0.999988
SPX500:1	0.011231	0.00409	2.75	0.006	0.003215	0.019247
ZN:1	1.708123	0.405528	4.21	0	0.913302	2.502944
constant	1433.749	399.9551	3.58	0	649.8511	2217.646

SPX500		*Coef.*	*Std. Err.*	*z*	*P>|z|*	*[95% Conf. Interval]*	
	JPY/AUD:1	-0.07622	0.019463	-3.92	0	0.114362	-0.03807
	AUS200:1	0.000201	0.000138	1.46	0.145	-6.9E-05	0.000471
	XS:1	-0.00017	3.81E-05	-4.47	0	-0.00025	-9.6E-05
	JGB:1	0.000504	0.000749	0.67	0.501	0.000965	0.001972
	JPN225:1	-2.2E-05	0.000033	-0.65	0.514	8.62E-05	4.31E-05
	SPX500:1	1.00055	0.000355	2818.46	0	0.999854	1.001246
	ZN:1	0.162499	0.035199	4.62	0	0.093509	0.231488
	constant	142.7208	34.71565	4.11	0	74.67938	210.7622

ZN		*Coef.*	*Std. Err.*	*z*	*P>|z|*	*[95% Conf. Interval]*	
	JPY/AUD:1	0.002073	0.000274	7.56	0	0.001536	0.00261
	AUS200:1	3.90E-07	1.94E-06	0.2	0.841	3.41E-06	4.19E-06
	XS:1	4.21E-06	5.36E-07	7.86	0	3.16E-06	5.27E-06
	JGB:1	-7.94E-06	1.05E-05	-0.75	0.452	2.86E-05	1.27E-05
	JPN225:1	-4.42E-07	4.64E-07	-0.95	0.341	1.35E-06	4.68E-07
	SPX500:1	-1.4E-05	5.00E-06	-2.89	0.004	2.42E-05	-4.65E-06
	ZN:1	0.995981	0.000496	2009.93	0	0.99501	0.996952
	constant	-3.62342	0.48872	-7.41	0	4.581291	-2.66554

A post-estimation test confirmed that the estimated VAR is stationary (all of the eigenvalues of the companion matrix lie outside of the unit circle). Additionally, a test for residual autocorrelation cannot reject the null hypothesis of no residual autocorrelation. A Jarque-Bera test indicates that the VAR disturbance terms are not normally distributed; this is good news because the LiNGAM algorithm relies on non-Gaussianity. Finally, both the LiNGAM and the FCI algorithms are applied to the estimates of the VAR innovation process; the resulting graphical structures are described below and shown in Figures 1 and 2.

LiNGAM Causal Search Algorithm Results

Graphviz software (*Graphviz*, 2019) was used to generate the graphs in Figures 1-3. The results from the LiNGAM algorithm, presented as two DAGs, are shown in Figure 1. The small DAG in the upper portion of Figure 1 contains the Australian and Japanese 10-year bonds and shows that the Australian bond influences the Japanese bond. Surprisingly, neither the Australian bond nor the Japanese bond have any connection to the JPY/AUD exchange rate.

The DAG in the bottom portion of Figure 1 shows that the JPY/AUD exchange rate is a source of information (all directed edges flowing out) for the three stock markets and the U.S. 10-year bond. The U.S. 10-year bond is an information sink (all directed edges flowing in) and is directly influenced by the JPY/AUD exchange rate and the Japanese and U.S. stock markets. The graph also shows that the Australian stock market leads the Japanese stock market, and both the Australian and Japanese stock markets have a causal influence on the U.S. stock market.

FIGURE 1
LiNGAM ALGORITHM OUTPUT

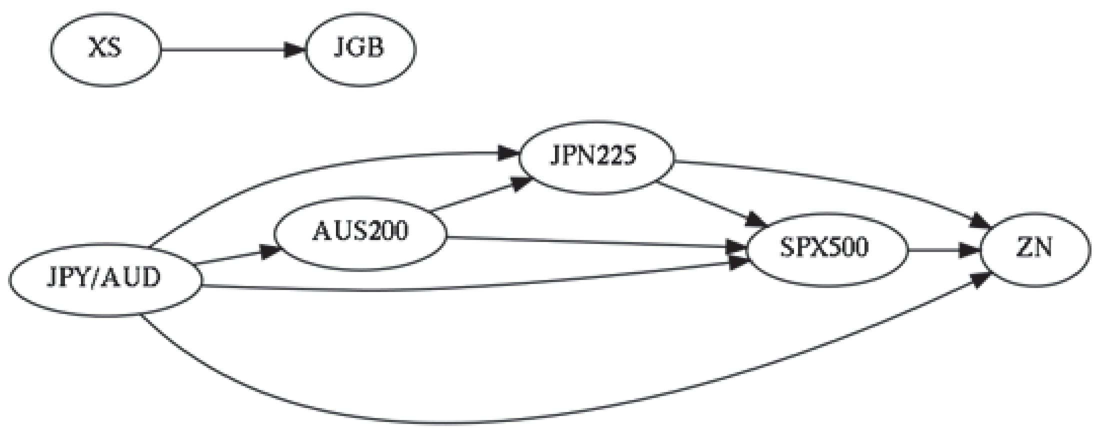

FCI Causal Search Algorithm Results

The results from the FCI algorithm are presented in Figure 2 as a partial ancestral graph. The FCI algorithm is unable to direct any edges in its graph, but the graphical skeleton is almost the same as that of the LiNGAM graph; the only difference is that the FCI graph does not contain the edge from the JPY/AUD exchange rate to the Australian stock market.

The lack of any directed edges in the FCI graph indicates, for example, that we do not know whether the Australian bond leads the Japanese bond or vice versa. We also cannot tell if there exists any possible common causes (because there are no edges with two arrow heads, e.g. A ↔ B). It is somewhat reassuring, however, that two dramatically different algorithms found all but one of the same edges in their graphical structures.

FIGURE 2
FCI ALGORITHM OUTPUT

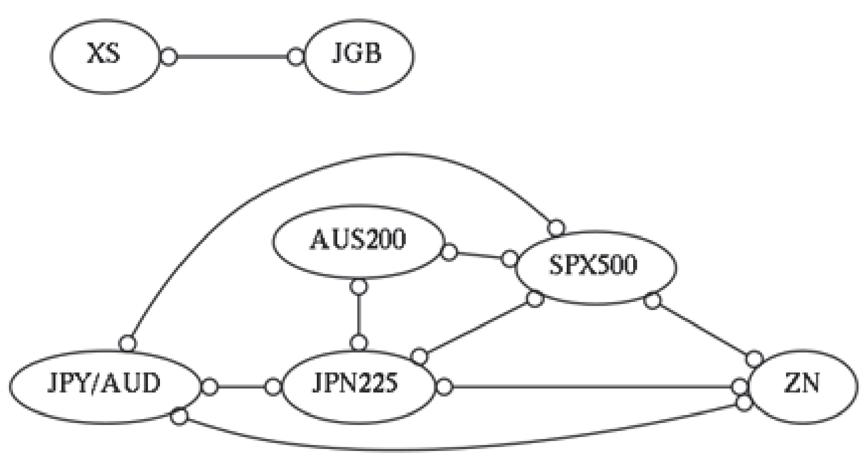

tsFCI Causal Search Algorithm Results

The results from the tsFCI algorithm are presented in Figure 3 as a partial ancestral graph. The tsFCI algorithm assumes the data is generated by a first-order time-invariant multivariate Markov process. The consequence of this assumption is that the algorithm's graphical output contains the variables at the current time (lag zero) and at one time period ago (lag one). In Figure 3 time flows from left to right so that the variable clusters on the left side of the figure with a colon and the number one (:1) appended to their names are the lag one variables while the variables on the right side of the figure are in the current period. Causal effects flow forward in time so that the lagged variables may have causal effects on the current variables but not vice versa.

The structure of the graphical skeleton in Figure 3 is similar to those of Figures 1 and 2 with the Australian and Japanese bond markets connected to each other but disconnected from the rest of the variables. Because there are open dots (uncertain endpoints) on the edges from the lag 1 variables to the lag 0 variables in the top part of Figure 3, it is unclear whether the lag 1 values of the Australian and Japanese bond markets have a direct causal influence on their lag 0 values or whether there is a latent common cause that affects both of their lag 1 and lag 0 values.

The direction of causal effects found by the tsFCI algorithm shown in Figure 3 is different from those found by LiNGAM shown in Figure 1 when looking within a particular slice of time. The tsFCI algorithm finds that the other variables have a causal influence on the JPY/AUD exchange rate so that it is now an information sink (within a time slice) instead of a source of information as shown in Figure 1. The tsFCI algorithm also finds that the U.S. stock market is an information source, but no definitive causal flows are found between any of the stock markets. There might exist latent common causes between the stock markets since there are open dots on the edges connecting them. The algorithm does find direct causal pathways from the U.S. stock market to the JPY/AUD exchange rate and to the U.S. 10-year note.

The edge with double arrowheads between the JPY/AUD exchange rate and the U.S. 10-year note indicates that a latent common cause (perhaps some other market?) may have a causal influence on both of them. The tsFCI algorithm also finds that the Japanese stock market has a causal influence on both the JPY/AUD exchange rate and the U.S. 10-year note.

There are three arrows directed from the lag 1 time slice to the lag 0 time slice. The lagged JPY/AUD exchange rate has a causal effect on the lag 0 values of both the JPY/AUD exchange rate and the Japanese stock market. The lagged value of the U.S. 10-year note has a causal effect only on its own lag 0 value.

FIGURE 3
TSFCI ALGORITHM OUTPUT

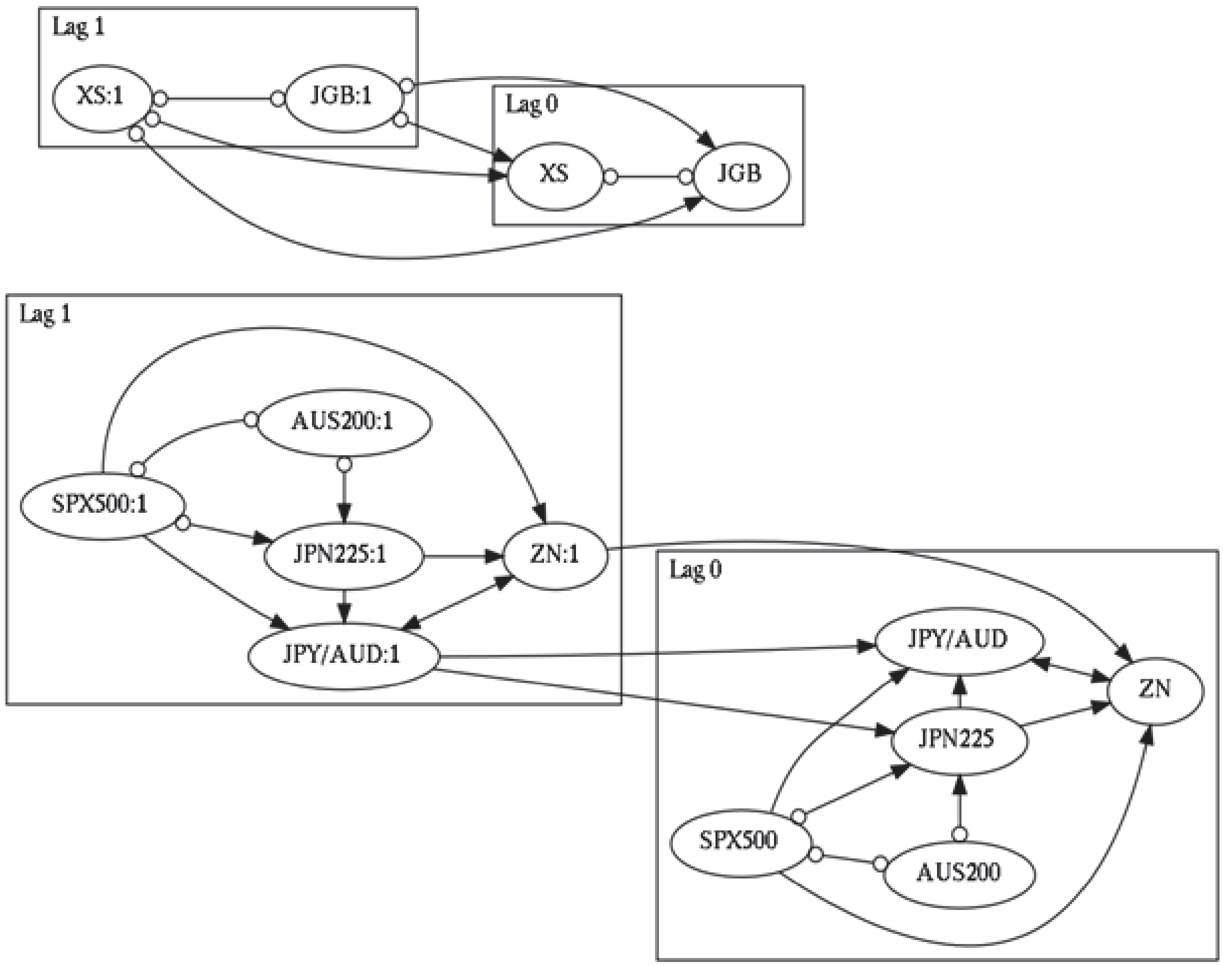

When compared to the output of the tsFCI algorithm, the results from the vector autoregression estimation shown in Table 1 contain many more linkages from past to present. There are a total of twenty-five variables with significant (p-value less than 0.01) coefficient estimates in the vector autoregression results versus only three edges from past to present in the tsFCI algorithm output. For example, there are four variables whose lagged values affect the current value of the S&P 500 index as shown in the second to the last panel in Table 1; this panel, labelled SP500, indicates that the lagged values of the JPY/AUD exchange rate, the Australian 10-year treasury bond, the U.S. stock market, and the U.S. 10-year note all have a significant effect (at the 0.01 significance level) on the current value of the U.S. stock market.

In summary, there are only indirect causal connections between the JPY/AUD exchange rate and the U.S. 10-year note and no connections between the JPY/AUD exchange rate and the Australian or Japanese bond markets. This finding provides evidence that the exchange rate market is not linked to the interest rate market as described by theory, and hence uncovered interest rate parity does not hold in the short run for higher frequency data. There are, however, direct causal connections between the JPY/AUD exchange rate and the U.S. and Japanese stock markets with the lagged value of the JPY/AUD directly affecting itself and the Japanese stock market at lag 0.

Extensions to This Research

Straightforward extensions to this research include adding one or more other variables such as additional stock or interest rate markets to the causal search. A more exotic study might include adding liquidity or volatility indicators such as the yield spread of Eurodollars over T-bills (TED spread) or the CBOE volatility index (VIX) to determine if they exert any influences on the other variables. Another useful goal would be to identify the latent common cause, found by the tsFCI algorithm, that lies between the JPY/AUD exchange rate and the U.S. 10-year note.

It is still an open question as to whether or not the models resulting from causal search algorithms are more useful than traditional time-series models. See (Deaton, 2018) as an example that addresses this issue by comparing the forecasting accuracy of causal search models to VAR and univariate autoregressive models. This topic needs further research.

CONCLUSIONS

All of the graphs in Figures 1-3 show that the Australian and Japanese bond markets are disconnected from the JPY/AUD exchange rate. This finding provides evidence that the theory of uncovered interest rate parity does not hold when using higher frequency (10 minute periodicity) data.

The JPY/AUD exchange rate is, however, connected to the U.S. stock market, the Japanese stock market, and the U.S. 10-year note in all of Figures 1-3. It appears that in this case the foreign exchange market moves with stock markets and the U.S. bond market rather than its own bond markets as theory would suggest. This finding is in line with six of the cited studies in the literature review section which also find that the carry trade is connected to one or more stock markets.

The LiNGAM algorithm finds that the JPY/AUD exchange rate is an information source and has a causal effect on the stock markets and the U.S. 10-year note. The FCI algorithm finds all but one of the same edges as LiNGAM but is unable to direct any edges.

The tsFCI algorithm finds that the U.S. stock market acts as an information source within a time-slice instead of the JPY/AUD as found by LiNGAM. The S&P 500 has a causal effect on the JPY/AUD exchange rate, Japanese stock market, and the U.S. 10-year note, and possibly the Australian stock market within a time-slice. Across time, the JPY/AUD exchange rate has a causal effect on itself and the Japanese stock market while the U.S. 10-year note influences only itself. And, finally, the tsFCI algorithm finds a latent common cause that affects both the JPY/AUD exchange rate and the U.S. 10-year note within a time-slice.

ACKNOWLEDGEMENT

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

REFERENCES

Bekaert, G., & Hodrick, R. J. (2012). *International financial management* (2nd ed). Boston: Pearson.

Carry on Speculating. (2007, February 22). The Economist. Retrieved from https://www.economist.com/finance-and-economics/2007/02/22/carry-on-speculating

Cheung, Y-L., Cheung, Y-W., & He, A. W. W. (2012). Yen carry trades and stock returns in target currency countries. *Japan and the World Economy*, 24(3), 174–183.

Christiansen, C., Ranaldo, A., & Söderlind, P. (2010). The Time-Varying Systematic Risk of Carry Trade Strategies. *Journal of Financial and Quantitative Analysis*, 46(4), 1107–1125.

CQG. (2019). Retrieved from https://www.cqg.com/

Deaton, B. (2018). Effects of the Swiss Franc/Euro Exchange Rate Floor on the Calibration of Probability Forecasts. *Forecasting*, 1(1), 2.

Entner, D., & Hoyer, P. O. (2010). On causal discovery from time series data using FCI. *Proceedings of the 5th European Workshop on Probabilistic Graphical Models*, 121–128.

Fung, H-G., Tse, Y., & Zhao, L. (2013). Are stock markets in Asia related to carry trade? *Pacific-Basin Finance Journal*, 25, 200–216.

Glymour, C., Scheines, R., Spirtes, P., & Ramsey, J. (2017). *The TETRAD Project*. Retrieved from http://www.phil.cmu.edu/projects/tetrad/

Graphviz. (2019). Retrieved from https://graphviz.org/

Nishigaki, H. (2007). Relationship between the yen carry trade and the related financial variables. *Economics Bulletin*, 2(13), 1–7.

Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge, U.K. ; New York: Cambridge University Press.

Ranaldo, A., & Söderlind, P. (2010). Safe Haven Currencies. *Review of Finance*, 10, 385–407.

Shimizu, S., Hoyer, P. O., Hyvärinen, A., & Kerminen, A. (2006). A Linear Non-Gaussian Acyclic Model for Causal Discovery. *Journal of Machine Learning Research*, (7), 2003–2030.

Sierra Chart. (2019). Retrieved from http://www.sierrachart.com

Spirtes, P., Glymour, C. N., & Scheines, R. (2000). *Causation, prediction, and search* (Second edition). Cambridge, Massachussetts London: The MIT Press.

Stata. (2019). (Version 14). StataCorp LLC. Retrieved from https://www.stata.com/

Tse, Y., & Zhao, L. (2012). The Relationship Between Currency Carry Trades and U.S. Stocks. *Journal of Futures Markets*, 32(3), 252–271.

Zhang, G., Yau, J., & Fung, H-G. (2010). Do credit default swaps predict currency values? *Applied Financial Economics*, 20(6), 439–458.