Journal of Banking & Finance 33 (2009) 578–582

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Is it the weather? Comment

Mark J. Kamstra [a], Lisa A. Kramer [b,*], Maurice D. Levi [c]

[a] Schulich School of Business, York University, 4700 Keele Street, Toronto, Ontario, Canada M3J 1P3
[b] Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, Ontario, Canada M5S 3E6
[c] Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, Canada V6T 1Z2

ARTICLE INFO

Article history:
Received 20 July 2008
Accepted 11 September 2008
Available online 24 September 2008

JEL classification:
G10
G11
G12

Keywords:
Seasonal affective disorder
SAD
Sell in May
Stock market cycles
Return seasonality

ABSTRACT

This comment discusses some errors in a recent paper by Jacobsen and Marquering [Jacobsen, B., Marquering, W., 2008. Is it the weather? Journal of Banking and Finance 32 (4), 526–540], in which the authors challenge our previous finding that stock market returns exhibit seasonal patterns consistent with the influence of seasonal affective disorder on investor risk aversion. We find that we cannot replicate the authors' findings, even after corresponding with them. Furthermore, we document several problems with their methodology, including misspecification of their economic model, misspecification of their econometric model, and use of inappropriate data. While we agree that seasonal affective disorder is not an explanation for all variation in equity markets, we do maintain that careful analysis leads to economically and statistically significant evidence of the effect we originally documented.

1. Introduction

Seasonality in stock returns has gone beyond a matter of speculating whether or not it exists to the question of which of a number of possible explanations may underlie it. This acceptance of return seasonality itself marks an important milestone, as does the fact that many of the principal explanations that are vying for general acceptance are based on human behavior.

Jacobsen and Marquering (2008; hereafter JM), set out along the path of challenging two behaviorally based explanations of seasonality, both related to investor mood, using a simple dummy variable that permits a shift in returns for half the year (November–April versus May–October). The explanations JM consider are temperature-induced mood shifts (Cao and Wei, 2005) and time-varying risk aversion induced by seasonal affective disorder (SAD) among investors (Kamstra et al., 2003). Among their core results is their conclusion that including a simple sell-in-May half-year dummy variable in the regression model eradicates the economic and statistical significance of the SAD effect. However, exploring the same data as JM we are unable to replicate their results, in spite of careful attempts and correspondence with the authors. We also note that JM make use of a misspecified economic model which mismeasures the SAD effect, they employ inappropriate return orthogonalizations which lead to an understatement of the SAD effect, and they use inappropriate data series for exploring the SAD effect: series that are too short and that are from countries with little or no seasonal variation in daylight, and where, therefore, one should not expect to find a SAD effect.

Before we consider the inappropriateness of the model and data used by JM it should be emphasized that a finding of a widespread sell-in-May effect would be interesting in its own right, but largely independent of the support we have found for time-varying risk aversion due to SAD. We see no reason why there may not also be a separate sell-in-May effect in financial markets from some reason that is yet to be identified.

In Section 2 we demonstrate that even when we replicate the estimation exercise reported by JM, we are unable to replicate their parameter estimates, finding statistically significant evidence of a SAD effect for many countries, in contrast to what they report. This occurs even though we use the JM model which is biased against finding a SAD effect. In Section 3, we enumerate additional problems with the JM methodology. For instance, we explain why the model they employ is misspecified, and why the majority of the indices they study are invalid for use in testing the SAD hypothesis in this context. We conclude in Section 4.

* Corresponding author. Tel.: +1 416 928 2496; fax: +1 416 971 3048.
E-mail addresses: mkamstra@yorku.ca (M.J. Kamstra), Lkramer@rotman.utoronto.ca (L.A. Kramer), maurice.levi@sauder.ubc.ca (M.D. Levi).

0378-4266/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2008.09.013

Table 1
SAD coefficient estimates

Country	Our SAD coefficient estimate (*t*-statistic)	JM's SAD coefficient estimate (*t*-statistic), extracted from Jacobsen and Marquering (2008)
USA	0.32	0.22
	(2.83)	(1.54)
UK	0.43	0.10
	(2.73)	(0.66)
Japan	0.60	0.32
	(2.02)	(1.10)
Sweden	0.43	0.29
	(3.14)	(1.98)

These regression results, using monthly data from January 1970 to May 2004, have country-specific MSCI value-weighted returns regressed on a constant, a January dummy variable, an NBER recession dummy variable, the return on the MSCI world index (orthogonalized with respect to the sell-in-May variable, consistent with the treatment JM apply but do not report in their paper), and a SAD variable which equals zero from April through September and equals the hours of night in excess of the annual average from October to March (with each country's length of night calculated at the latitude of the country's major exchange). The *t*-tests for our results are based on MacKinnon and White (1985) jackknife heteroskedasticity-robust standard errors. The data we employ are US-denominated, though we find similar results using local-denominated returns. We do not know whether JM employ US- or local-denominated returns.

2. Replication of the model Jacobsen and Marquering explore

We downloaded the MSCI value-weighted country index data for the 48 countries JM investigate, and have attempted to replicate their findings. We have been able to find qualitatively similar results to those they report for their sell-in-May dummy variable. However, we have been unable to replicate other important aspects of their reported results, even after corresponding with the authors.[1] Most importantly from our perspective, we have not been able to replicate, even qualitatively, their results for the SAD length-of-night variable, even when using the same data range they report, the same monthly frequency they employ, the same variables they use (including some that we do not believe belong in the model, as we explain in footnote 1), and the same misspecified model (inappropriately excluding the fall dummy variable, a point we address more fully below).

In Table 1 we provide some of the starkest differences we found between JM's reported results and the ones we produced by exactly replicating their steps. For the small set of series we present, and in general for the most interesting series (by market capitalization), we find much stronger results for the SAD variable than reported by JM. That is, overall, we find that JM report systematically smaller SAD coefficient estimates and smaller *t*-statistics than we were able to replicate. In sum, even if we estimate JM's model (a model which is not correctly specified to test for the SAD effect), we often find more significant coefficients on the SAD length-of-night variable than they report.

[1] In Ben Jacobsen's response to our query, he alerted us to some estimation details that were not provided in Jacobsen and Marquering (2008). These bear reporting here, to aid other researchers attempting to replicate their findings. Their Table 1 results do not use the MSCI world return as a regressor, inconsistent with the notes to the table. Instead, the authors first orthogonalize the MSCI world return with respect to the sell-in-May dummy variable, and then in all the regressions reported in their Table 1, use the residuals from this orthogonalization regression as the explanatory variable referred to as the return on the MSCI world index. This treatment advantages the sell-in-May variable relative to other seasonal variables in the regression, making it more likely that the sell-in-May variable will be found to be economically and statistically significant and less likely that others will. We elaborate on this point more fully below.

3. Additional problems

In addition to the problems mentioned above, there are several others which we detail in this section.

The authors assert (in footnote 15, on page 531) that their results for our SAD model are qualitatively similar whether or not they include the fall dummy variable of the original Kamstra et al. (2003) model specification. In drawing comparisons, it would have been helpful if they had presented the results based on our original model (which incorporates a fall dummy variable). This would facilitate direct comparison with the published results and allow readers to determine the source of any differences in results.

Some background on the SAD effect may help the reader understand why we included a fall dummy variable in our original study.[2] As we stated in our 2003 paper, the psychology literature has established that SAD induces depression, and separately, that depression is associated with heightened aversion to risk, including risk of a financial nature. All of the clinical evidence on the incidence of SAD shows that onset of the condition tends to occur in early fall, as the amount of daylight diminishes (literally as the amount of time between sunrise and sunset shortens), and recovery occurs as the length of daylight expands in the new year.[3] The proportion of the population suffering from this affective disorder rises during the fall and then declines with the approach of spring. SAD incidence is directly related to willingness to bear risk. The implication of investors' risk aversion increasing during fall and alleviating during winter or spring is that, all else held constant, stock returns should be lower in the fall (as SAD-affected investors shun risky securities) and higher in the new year (as investors recovering from SAD resume their risky holdings). JM point to an observation raised by Kelly and Meschke (2007) that "depression peaks due to SAD did not occur during the fall but during the period December–February (page 529, JM)." The authors are confusing *flow* and *stock* concepts here. Equity returns, an income *flow*, respond to the *flow* of SAD-affected investors, not the *stock* of SAD-affected investors. It is the *flow* (the onset and then the recovery) that we hypothesize moves markets, with newly affected investors rebalancing their portfolios to reflect their changing risk tolerance. Testing the impact of the *stock* of SAD-affected investors on returns necessarily mixes dimensions and hence misspecifies our model. It is not how many people that have SAD that matters; it is how many more or less SAD-affected people are rebalancing their portfolios that matters. That is, it is the flow of new or recovering SAD-affected people that matters. Thus it is incidental whether or not SAD patients "feel worst" during the December through February period. The pertinent issue is the *timing of onset* of and *recovery* from SAD. September and October are actually the months when the highest proportion of individuals *start* suffering from SAD (see Kamstra et al., 2008a), so if people start rearranging their portfolios when they first become risk averse, those months should arguably be the times when we see the biggest negative impact on equity returns due to SAD. The mirror image would be expected to occur in the winter. Some people start recovering in January, but the peak month for recovery is March. So we should see positive effects in equity returns as early as January, but the peak effect should not take place until March. The JM model explicitly re-

[2] As we explain below, we have since found that use of a single alternative variable (based on the clinical incidence of SAD in populations known to suffer from the condition) allows one to avoid the use of a fall dummy variable (see Kamstra et al., 2008a). However, when using the 'length-of-day' variable from our original study, we do advocate inclusion of a fall dummy variable.

[3] We note in passing that SAD does not seem to be related to the amount of sunlight versus cloud cover, cloud cover being a regional phenomenon which can of course differ across cities and across days within a season. Nor does SAD seem to be related to other aspects of weather which vary from day to day and across regions, such as precipitation or temperature. We refer the interested reader to Kamstra et al. (2003) for the associated citations to the medical literature.

stricts returns due to SAD from differing across the fall and winter seasons as it excludes the fall dummy variable, and therefore does not permit testing of the Kamstra et al. (2003) SAD hypothesis. When we advocated in our 2003 paper for the inclusion of a fall dummy variable in the model specification, we aimed explicitly to allow for the possibility that returns should be lower in the fall and higher in the winter under the SAD hypothesis.

Turning to other issues, while it is noteworthy for JM to have collected data for 48 different countries, we do not believe it is advisable to test for the influence of SAD on financial markets in most of those countries. First, in the context of testing for seasonal effects, it is sensible to avoid using data from countries that have experienced hyperinflation (even if the returns are measured in US dollars), because hyperinflation distorts virtually every aspect of an economy's financial system. JM consider countries which have experienced hyperinflation over the period they study, such as Brazil and Argentina. Second, it is best to avoid data from markets that are closely linked to particular commodities, since commodity prices frequently exhibit strong seasonal variations of their own: many of the exchanges JM study are dominated by firms in commodity-intensive sectors, such as Venezuela where oil accounts for roughly a third of the country's GDP. Third, it is not theoretically justifiable to test for SAD using data from equatorial countries, since there is very little seasonal variation in daylight in such countries. For example, in countries within 20° of the equator, of which JM include 10 (of the 48 countries they study in total), the number of hours of daylight varies only ±1 hour around its annual average. Investors are highly unlikely to experience time-varying risk aversion due to SAD in such equatorial locations. Fourth, even southern hemisphere exchanges are somewhat problematic to the extent that international equity markets are integrated and northern hemisphere investors (who comprise the bulk of international wealth and investors) dominate mature markets like New Zealand and Australia. (This likely explains the somewhat weaker results we found in the southern hemisphere countries we originally considered.) Fifth, when studying a phenomenon such as SAD which occurs at an annual periodicity, it is important to use long time series. Yet for many of the countries JM study, as little as 10 years of data is used, and in no cases are data preceding 1970 employed. For all of the above reasons, we limited our original study to broadly-based equity markets in countries with modest inflation and where long series of reliable daily data are available.

On the topic of returns frequency, JM employ monthly rather than daily data. They describe daily data as "noisy," but it is actually helpful to consider daily data when testing a hypothesis which has implications for data at a daily frequency, as does the SAD hypothesis. (Daylight changes not just from month to month, but also from day to day, and so use of daily data allows for more powerful tests of the SAD hypothesis.) Subtleties of the daily changes in daylight are lost when testing the SAD hypothesis at the monthly frequency, biasing tests against the SAD hypothesis. Additionally, it can be helpful to consider daily rather than monthly data when trying to disentangle the separate influences of factors such as SAD and temperature, which themselves have similar time-series properties and which distinguish themselves only in subtle ways. In our own research we often present additional results based on monthly data as a robustness check, as in Garrett et al. (2005), but we do not believe it wise to consider monthly returns exclusively.

There are several problems with JM's choice of explanatory variables for their model. Difficulties arise from their decisions to ignore possible effects due to tax-loss selling (tax years commence in different months of the year for the various countries they consider) and to include an additional January dummy variable for *all* of the countries they consider (including those for which the tax year does not begin in January). The month of January comes at a critical time of the year for the time path of SAD, and by restricting the seasonal behavior of returns in that month to be attributed solely to a January dummy variable, even when the tax year does not begin in that month, the estimation biases the outcome against the SAD hypothesis. Furthermore, JM's Table 3 (page 538) includes a dummy for the October effect. We do not know of an October equity return anomaly. A literature survey reveals one article which documents an October anomaly in returns to *defaulted bonds* (see Ward and Huffman, 1997), and another article by Kryzanowski and Zhang (1992) which reports an *in*significant October equity return with Canadian data. However, we can find no work that documents a systematic October effect in equity returns. Including *ad hoc* dummy variables for various months of the year when using monthly data to test for an annual seasonality can badly bias tests for seasonality. This practice is particularly problematic when there is no literature documenting the existence of the anomaly for which the dummy variable is intended to control. Related to this point is the inclusion of a dummy variable for NBER-dated recessions. The NBER recession variable is known to exhibit an *ex post* bias: in some cases recessions are not labeled as such by the NBER until years after the event. The Stock and Watson (1989) real-time probability-of-recession variable might be a suitable replacement (the Federal Reserve Bank of Chicago maintains an up-to-date probability-of-recession variable), but even the use of that variable would be unusual in the context of capital market regressions.

As we mention in footnote 1, the authors employ as one of their control variables the MSCI world-index return orthogonalized with respect to the sell-in-May variable. In this context, the use of a market return variable as an explanatory variable is problematic. The overall market return itself exhibits seasonal patterns, including the SAD effect and the sell-in-May effect. Thus if one were to use the (unorthogonalized) world-index return, one would be unable to find separate evidence of the SAD effect or the sell-in-May effect. (Those effects would be subsumed by the world-index return.) This is likely why the authors elected to orthogonalize the world-index return with respect to the sell-in-May variable. However, their choice to orthogonalize with respect to *only* the sell-in-May effect means that they will be unable to find separate evidence for the presence of the SAD effect (except for any portion of the SAD effect which happens to be correlated with the sell-in-May variable); they will, however, still be able to find evidence of the sell-in-May effect, having previously orthogonalized the world-index return with respect to this variable.

Further on the issue of using a world-index return variable, when Fama and French (1993) identify common risk factors in the time-series returns to stocks and bonds, they find that the shared impact of these factors across stock and bond returns appears to come in through the excess market return, which is itself influenced by all the factors. To distinguish the roles of the bond and equity factors, Fama and French (1993) orthogonalize the excess market return with respect to these factors, and use this orthogonalized variable in place of the excess return on the overall market. Then, when they run their regressions (including as a regressor the excess market return that has been orthogonalized with respect to the bond and equity factors, and including the bond and equity factors themselves as regressors), the importance of the bond and equity factors can be accurately evaluated on the basis of their coefficient estimates. Had Fama and French not first orthogonalized the market return with respect to these variables, they would not have found evidence that the bond factors influence the bond portfolio returns. Thus, when including a market return in their regression (something which implies a capital asset pricing framework) JM should orthogonalize the market return with respect to all of the seasonal variables they seek to test, and they should consider *excess* market returns instead of raw returns. We employ such an orthogonalization technique in Kamstra et al. (2008a) and find that doing so leads to even stronger support for the SAD hypothesis. The current specification employed by JM, however, understates the impact of SAD by including a variable correlated with SAD (the world return) and not first orthogonalizing it with respect to SAD.

Turning to issues relating to the theory underlying the SAD hypothesis, JM speculate that investors who work indoors may be immune to the effects of environmental variables. While this possibility may be casually intuitive, when it comes to SAD, the medical literature indicates that for affected individuals who work indoors, the impact of the reduced daylight through the fall and winter is at least equivalent to that for SAD sufferers who work outdoors. Indeed, it may even be *more* severe (for clinical evidence, see Wirz-Justice et al., 1992 and Magnusson and Stefansson, 1993). Further issues relating to the literature on SAD and/or depression include several references to a marketing study by Parker and Tavassoli (2000). The authors note that Parker and Tavassoli "argue that not depressed people but [rather] people in positive moods seem to become more risk averse." In fact, Parker and Tavassoli's article does not comment on the risk aversion of depressed people and does not dispute the well-established finding that depressed individuals are more risk averse. The authors also note that Parker and Tavassoli "indicate that lack of sunlight might arouse risk-taking behavior," but Parker and Tavassoli's paper, titled "Homeostasis and consumer behavior across cultures," was written in the context of attempting to predict, for instance, how consumers might be more likely to buy fattier and sweeter candy bars in colder regions. Additionally, Parker and Tavassoli's hypothesis is *speculative* (without any clinical or empirical support), and furthermore, Parker and Tavassoli make *no* reference to financial risk aversion.

We recently introduced (see Kamstra et al., 2008a) an alternate specification for modeling SAD which addresses many of the concerns raised by the authors and by Kelly and Meschke (2007). The new specification employs a variable that is based directly on the *clinical incidence of SAD symptoms among individuals who suffer from the condition*, avoiding use of both the ad hoc fall dummy variable and the "complicated trigonometric formulas" (a sine wave that approximates the length of night) mentioned by JM. We find that the new variable based on the clinical incidence of SAD is at least as effective in explaining seasonal patterns in equity returns as the two-variable specification in our 2003 paper. The new variable is available on the web, both at daily and monthly frequencies, at www.lisakramer.com/data.html. We recommend its use in place of the length-of-night and fall-dummy-variable specification when testing for the influence of SAD on financial markets.

We should also note, for the sake of any researcher attempting to undertake tests of SAD, that there are some errors in the latitude information provided in JM's Table 1 (page 532). Latitudes are conventionally expressed in degrees and minutes, and there are sixty minutes in each degree. In JM's Table 1, the latitude figures for many countries are expressed with more than 60 minutes.

4. Conclusions

There are dimensions of alternative explanations of seasonality that might be able to distinguish between variables as closely connected as hours of daylight and temperature, or what happens after May or leading up to Halloween. For example, a SAD-based explanation working through time-varying risk aversion would suggest an opposing seasonal pattern in low-risk fixed income securities relative to the pattern Kamstra et al. (2003) show in equity returns. There would be higher returns in Treasury bonds during periods when SAD-affected investors are shunning risky securities and lower Treasury-bond returns during periods when SAD-affected investors are willing to tolerate more risk in their portfolios. Such an opposing seasonal pattern in Treasury-bond returns has indeed been identified by Kamstra et al. (2008a). Further corroborating evidence has been provided by Kamstra et al. (2008b) who find seasonal patterns consistent with SAD in funds flowing between safe and risky categories of mutual funds during the year, and by DeGennaro et al., (2008) who find SAD-consistent evidence that market makers exhibit seasonal variation in the spreads between their bid and ask quotes. Related work supporting the notion of SAD-induced time-varying risk aversion has also been shown by Dolvin and Pyles (2007) who document SAD in the returns to stocks that have undergone an initial public offering, and Kaplanski and Levy (2008) who document a relationship between SAD and the Chicago Board Options Exchange Volatilty Index (VIX) which is known also as the "Fear Index." Furthermore, Dowling and Lucey (2008) study equity returns from 37 countries, employing much of the same data as JM, and find strong evidence in support of the SAD effect.

Certainly, the SAD effect does not explain everything. In our original paper we state clearly that the SAD effect is more likely to be present in large non-equatorial countries with broad-based, diversified economies. Our careful review of the paper by Jacobsen and Marquering leaves us unshaken in this conclusion. We remain convinced that a properly specified model applied to a reasonably long time series of daily and/or monthly data for a well-motivated set of indices yields strong evidence in support of an economically meaningful and statistically significant SAD effect related to variation in investor risk aversion through the year.

Acknowledgements

The authors are grateful to the Social Sciences and Humanities Research Council of Canada for financial support. We thank the editor, Ike Mathur, and Anthony Saunders for helpful suggestions. Any errors are our own.

References

Cao, M., Wei, J., 2005. Stock market returns: a note on temperature anomaly. Journal of Banking and Finance 29 (6), 1559–1573.
DeGennaro, R., Kamstra, M.J., Kramer, L.A., 2008. Does risk aversion vary during the year? Evidence from bid-ask spreads. Working Paper, University of Toronto.
Dolvin, S.D., Pyles, M.K., 2007. Seasonal affective disorder and the pricing of IPOs. Review of Accounting and Finance 6 (2), 214–228.
Dowling, M., Lucey, B.M., 2008. Robust global mood influences in equity pricing. Journal of Multinational Financial Management 18, 145–164.
Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33 (1), 3–56.
Garrett, I., Kamstra, M.J., Kramer, L.A., 2005. Winter blues and time variation in the price of risk. Journal of Empirical Finance 12 (2), 291–316.
Jacobsen, B., Marquering, W., 2008. Is it the weather? Journal of Banking and Finance 32 (4), 526–540.
Kamstra, M.J., Kramer, L.A., Levi, M.D., 2003. Winter blues: a SAD stock market cycle. American Economic Review 93 (1), 324–343.
Kamstra, M.J., Kramer, L.A., Levi, M.D., 2008a. Opposing seasonalities in Treasury versus equity returns. Working Paper, University of Toronto.
Kamstra, M.J., Kramer, L.A., Levi, M.D., Wermers, R., 2008b. Seasonal asset allocation: Evidence from mutual fund flows. Working Paper, University of Toronto.
Kaplanski, G., Levy, H., 2008. Seasonal affective disorder (SAD) and perceived market risk. Working Paper, Hebrew University of Jerusalem.
Kelly, P.J., Meschke, F., 2007. Sentiment and stock returns: the SAD anomaly revisited. Working Paper, University of South Florida.
Kryzanowski, L., Zhang, H., 1992. Economic forces and seasonality in security returns. Review of Quantitative Finance and Accounting 2, 227–244.
MacKinnon, J.G., White, H., 1985. Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. Journal of Econometrics 29 (3), 305–325.

Magnusson, A., Stefansson, J.G., 1993. Prevalence of seasonal affective disorder in Iceland. Archives of General Psychiatry 50, 941–951.

Parker, P.M., Tavassoli, N.T., 2000. Homeostatis and consumer behavior across cultures. International Journal of Research in Marketing 17, 33–53.

Stock, J.H., Watson, M.W., 1989. New indexes of coincident and leading economic indicators. In: Blanchard, O.J., Fischer, S. (Eds.), NBER Macroeconomics Annual. MIT Press, Cambridge.

Ward, D.J., Huffman, S.P., 1997. Seasonality in the returns of defaulted bonds: the January and October effects. Quarterly Journal of Business and Economics 36, 3–10.

Wirz-Justice, A., Kräuchi, K., Graw, P., Schulman, J., Wirz, H., 1992. Seasonality in Switzerland: an epidemiological survey. Society for Light Treatment and Biological Rhythms Abstracts 4, 33.