Acquisition of information and share prices: An empirical investigation of cognitive dissonance*

Elena Argentesi† **Helmut Lütkepohl**‡

Massimo Motta§

This draft: 17 November 2007
First draft: 23 March 2005

Abstract

This paper deals with the determinants of agents' acquisition of information. Our econometric evidence shows that the general index of Italian share-prices and the series of Italy's financial newspaper sales are cointegrated, and the former series Granger-causes the latter, thereby giving support to the cognitive dissonance hypothesis: (non-professional) agents tend to buy the newspaper when share prices are high and not to buy it when share prices are low. Instead, we do not find support for the hypothesis that the agents acquire information in order to trade in the stock-market: we find no relationship between quantities exchanged in the market and newspaper sales, nor between stock market volatility and newspaper sales.

*We are grateful to Armando Carcaterra, Claudio Mazzoli (Brainpower) and Hammad Nasar for their help with Italian stock-market data, to Sibilla Guzzetti (ADS) for help with the newspapers sales data, and to Andrea Beltratti, Giuseppe Bertola, Laura Bottazzi, Antonio Cabrales, Juan Carrillo, Stefano DellaVigna, Luigi Guiso, Elisabetta Iossa, Tullio Jappelli, Bruno Jullien, Gregor Langus, Annamaria Lusardi, Marco Pagano, Ludovic Renou, Riccardo Rovelli, Karl Schlag, and seminar participants at the EUI (Florence), Bologna, WZB (Berlin), and ESEM 2006 (Vienna) for useful discussions.

†University of Bologna. Address: Department of Economics, Piazza Scaravilli 2, I-40126 Bologna. E-mail: elena.argentesi@unibo.it.

‡European University Institute, Florence. Address: EUI, Department of Economics, Via della Piazzuola 43, I-50133 Firenze.

§European University Institute, University of Bologna and CEPR. Address: EUI, Department of Economics, Via della Piazzuola 43, I-50133 Firenze.

1 Introduction

In 'standard' economic theory, an agent will normally be better off by having more information, if the latter was free.[1] This is because utility depends on outcomes, and information (if it has any relevance) should help an agent take better decisions, in turn improving outcomes. In recent years, however, a growing literature has integrated psychology into economics, suggesting various reasons why agents might want not to acquire available information. Different ways to model this phenomenon have been proposed, including strategic behavior by agents, as well as the incorporation of beliefs in the utility function of individuals (see below for references).

To investigate empirically whether information acquisition is driven by psychological considerations, we look at the relationship between non-professional investors' acquisition of information about financial markets - in the particular form of their purchase of Italy's main financial newspaper, *Il Sole 24Ore* - and data on the evolution of the Italian stock exchange market.

Our hypothesis is that this relationship is consistent with the theory of cognitive dissonance (Festinger, 1957), very influential in social psychology, and supported by a number of anecdotal and experimental findings, in very different domains and contexts.[2] This theory postulates that two cognitions (or elements of knowledge) are dissonant if the opposite of one cognition follows from the other. Dissonance makes an agent uncomfortable, and in order to reduce it the agent may either avoid any information likely to create dissonance or process the available information so as to reduce it. In our context, an agent who learns that the price of the shares she holds decreases will experience dissonance: the cognition that a share is 'doing badly' is dissonant with the cognition that she holds that share in her portfolio. In order to reduce dissonance the agent can

[1] The idea that in single-person decision problems agents prefer more information to less was formalized by Blackwell (1951, 1953). This may not be the case in a multiple-agent setting (see Hirshleifer, 1971). In game theory, there are many situations in which players are worse off with more information (see also Osborne, 2004 p. 283 for examples).

[2] Two economics papers which report evidence on cognitive dissonance and the related phenomenon of confirmatory bias are Rabin (1998) and Yariv (2002).

keep her behavior unchanged (i.e., continue to hold the share) while eliminating the dissonant knowledge by ignoring information about the share price.[3] Therefore, at an aggregate level this theory suggests that agents acquire information when the share-price index increases (that is, they buy the financial newspaper when they expect to see that the particular shares they hold are doing well), but prefer to ignore information when the share-price index decreases (that is, they do not buy the newspaper when they expect to learn from it that their shares are not doing well). We assume that expectations are correct on average, as the agent is exposed to some rumour about the general price level, but it is only after buying a newspaper that she will have precise information on her shares.[4]

By using cointegration techniques, we find that our data (we have monthly observations from 1978 to 2003) lend support to the cognitive dissonance hypothesis. The share-price series and the series of the financial newspaper sales are cointegrated (i.e., they move together), and the former causes (in the sense of Granger, 1969) the latter.[5]

We also analyse whether more "standard" hypotheses about information acquisition are consistent with the data. For instance, agents may buy newspapers to acquire information about the stock market, in order to improve their

[3] Another way in which dissonance may be reduced if negative news on the share-price index appear, could be to sell her shares. However, this remedy to dissonance is certainly more costly (selling shares would entail a transaction cost, plus the agent should take another decision on how to invest the money realized from the sale) than simply ignoring information (note also that information here is costly: not buying the newspaper entails a saving). Furthermore, this behaviour would also contrast with *prospect theory* (itself supported by some empirical evidence): people tend to hold on their shares when they are doing badly (with respect to some benchmark). This behaviour, known in the finance literature as the "disposition effect", is a well documented empirical regularity (see Odean, 1998, and the subsequent literature).

[4] Another way to formalise this idea is that, if share prices are correlated over time (as it is argued by the financial literature on 'mean reversion', see for instance Fama and French, 1988), then the agent may buy the newspaper for one or more days, but after having observed a drop in the asset prices would not buy it any longer for some time after, since she expects it would still report low prices.

[5] We posit here (and find) a casual link from asset prices to newspaper sales. Dyck and Zingales (2003) analyze a somewhat opposite causal link, by asking whether media coverage affect asset prices. However, what could affect stock-exchange performance are particular news about companies and sectors (for instance, expectation and announcements of earnings, news about demand and cost evolution). It is difficult to find a reason why the *number* of copies sold of a financial newspaper should affect positively the stockmarket prices. There may be of course events that both increase sales of the newspaper(s) and affect the stock-exchange prices, but one can reasonably expect that some of these events will negatively affect prices, while others will positively affect them.

trading in the market, or reoptimise their portfolio (the newspaper may report detailed information about share prices, useful to better calculate own portfolio allocation). According to this view, one should expect to find that the financial newspaper sales increase with the volumes traded in the market.[6,7]

Another prediction of the rational model is that the proportion of informed individuals increases with price noise (see for example Grossman and Stiglitz, 1980). This is due to the fact that the higher the level of noise, the less informative the price system is, and therefore the more valuable information is to traders.

However, we find that the volumes traded in the stock exchange are not cointegrated with the financial newspaper sales, nor is there any evidence of a causal relation between stock market volatility and newspaper sales. Therefore our analysis does not lend support to these two particular "rational" explanations of information acquisition in financial markets.

We also perform a series of robustness checks (see section 3 for a discussion). Firsly, we find no evidence of the existence of reverse causality relations between the variables. Further, we try to rule out some confounding factors which could possibly drive the relationship between the stock market index and the financial newspaper sales. In particular, it may be that when the stock market performs well, people increase their consumption of any type of newspaper. We perform a cointegration analysis between the stock price index and the sales of sports newspapers (which are probably the class of newspapers more distant from the financial ones). We find no evidence of cointegration between these

[6] The fact that more infomation is collected by investors does not necessarily imply that more trade will follow (for instance, because information may just suggest that it is optimal not to trade). However, under a neoclassical hypothesis one should presume that if the volume of trade in the market increases, investors should have previously collected all the available information.

[7] The idea that more informed individuals trade more is an implication of rational models on the determinants of investment in financial information. For example, in Peress (2004)'s model, information is more valuable to agents with a riskier portfolio (who are also wealthier investors). These agents therefore acquire more information, which increases the precision of their signal. In turn, a higher precision induces more informed agents to hold more stocks. Guiso and Jappelli (2005) propose an alternative model which yields the same prediction as the rational model, i.e. more informed agents trade more. However, in their model, the driving force is a behavioral motivation: agents are overconfident about the quality of the information, and therefore trade more in response to the information collected.

variables, which suggests that, unlike the consumption of financial papers, the consumption of sports papers does not track stock market performance. Finally, we show results of a causality analysis between the Italian industrial production (used as a proxy for economic activity) and Sole 24Ore sales. The two series are not cointegrated nor there is trace of a causal relation (neither Granger nor instantaneous) between them.

In order to provide additional evidence on the effect of interest, we replicated our analysis using data from the United Kingdom (see section 4). Our results show that there is evidence of a cointegration relation between the UK stock market index and the Financial Times circulation, and that the former Granger-causes the latter.[8]

The advantages of using Italian data are two-fold. First of all, data for newspaper sales are available on a longer sample period (in the UK data there are several missing observations before 1985, therefore we only use the sample period without missing data). Secondly, the slower diffusion of the internet in Italy with respect to other countries allows us to credibly consider the financial newspaper as the main source of information on financial markets during most of our sample period.

Our paper is related to the recent and fertile literature on economics and psychology.[9] More particularly, a number of distinct models have been developed which are able to explain biases in the acquisition of information.[10]

Among the more recent contributions,[11] Rabin and Schrag (1999) explain the existence of distortions in agents' information through cognitive mistakes, and Carrillo and Mariotti (2000) explain 'anomalous' attitudes to information through strategic decisions of agents, who choose to be ignorant in order to

[8] We were unable to replicate the analysis on US data since the US Audit Bureau of Circulation does not provide monthly data on newspaper sales (they only have quarterly data, a frequency which would considerably reduce the power of our econometric analysis).

[9] See for instance Rabin (1998, 2002), Brocas and Carrillo (2003, 2004), and Camerer, Loewenstein and Rabin (2004).

[10] It is not surprising that the same phenomenon is explained by distinct approaches, since the theory of cognitive dissonance has given rise to a large variety of interpretations and applications. See Harmon-Jones and Mills (1999) for a review of the social psychology literature.

[11] See Akerlof and Dickens (1982) for the first formalization of cognitive dissonance within an economic model.

discipline their future behavior.[12]

A third approach explains cognitive dissonance by assuming that the agent's beliefs enter directly her expected-utility function. This approach has been pioneered by Akerlof and Dickens (1982) and recent contributions include Köszegi (2001), Eliaz and Spiegler (2003), Yariv (2005).[13] In this approach, information can be used by agents to improve their decisions, but it can also affect their beliefs. An agent who maximizes a standard expected-utility function would not refuse to have free information because this would allow her to take 'better' decisions, but an agent whose beliefs enter her expected-utility function may decide to ignore information (or to re-interpret it) so as to preserve her beliefs. 'Anomalous' behavior with respect to information follows from the type of beliefs that the agent has. For instance, if the agent has a preference for consistent beliefs, her utility increases when her beliefs are confirmed, and decreases when they are invalidated: the agent may want to actively acquire information of the former type, and to ignore (or manipulate) information which leads to the latter situation. Or, if the agent's utility increases with the belief the agent holds about herself, the agent may want to ignore any information which would lead her to revise downwards the judgment of her abilities.

Although we do not venture into a theoretical model of our findings, we speculate that the third approach might naturally lead to agents' behavior consistent with our empirical findings. Suppose that an agent's expected utility includes not only the performance of the assets she holds, but also her beliefs on her abilities as an investor. Then, our agent should be eager to acquire positive news about the performance of her assets and would instead prefer not to see the negative news. In other words, she would buy the financial newspaper in times of high share prices and not buy it in times of low share prices.

In this stream of the literature, Karlsson et al. (2005) present a model of belief manipulation. When facing a changing environment, agents choose between two psychological states: they can be either attentive and actively

[12] See also Benabou and Tirole (2002).

[13] In all of these papers, beliefs enter directly the utility function of the agent, but in some works beliefs are treated as a choice variable, while others treat them as parameters.

seek information, or inattentive and avoid information. The authors find that for some parameter values their model gives rise to what they term 'ostrich effect' (and we call cognitive dissonance): in 'bad times', individuals choose to be inattentive (and put their heads in the sand like ostriches), while in 'good times' they choose to be attentive.

What makes their paper similar to ours is that they also investigate this question empirically by looking at share prices data. They find evidence that the aggregate number of daily logins to investors' online accounts is positively related to the stock exchange prices, implying that investors are more likely to check the value of their portfolio when the market is up. This finding therefore provides additional evidence in favor of the 'cognitive dissonance' hypothesis on a different dataset and with a different methodology than ours.

There are two main differences between Karlsson et al. (2005) and the present paper. Firstly, and more importantly, we do not limit ourselves to investigating whether the data support the 'cognitive dissonance' (or 'ostrich behavior') hypothesis, but we also investigate competing hypotheses, according to which information is acquired in order to improve decision-making. Secondly, we make use of different econometric methods. They have daily data for a relatively short period of time, and they limit themselves to simply regress the aggregate number of daily logins on the relevant share price indices; instead, we have monthly (instead of daily) observations but much longer (26 years) series of data, and we use more sophisticated econometric techniques which also allow us to investigate the (Granger-)causality link between the variables at issue.

The remainder of the paper is as follows. Section 2 describes the data; section 3 presents our econometric methodology and results, and carries out a series of robustness checks; section 4 considers UK data and section 5 concludes.

2 Discussion of the data

We use average monthly values from 1978 to 2003 for the Datastream price index of the Italian stock market. This index is built by taking the first 90% of all the companies quoted on the Italian stock exchange taken in decreasing order of capitalization, disregarding in this way the small companies. The trading volumes corresponding to the Datastream stock price index are only available from 1986 to date (we use total monthly volumes).

As to the sources of information on the financial market, we use data on the sales of *Il Sole 24Ore*, which is by far the largest financial (daily) newspaper.[14] The Italian market for financial newspapers is essentially composed of three titles: *Il Sole 24Ore*, *Italia Oggi*, and *MF*.[15] However, *Italia Oggi* and *MF* have a much lower circulation with respect to *Il Sole 24Ore*, which has historically always been *the* Italian financial newspaper and accounts for over 90% of this market.[16] Therefore we only consider the sales of *Il Sole 24Ore* for the purpose of our analysis. In particular, consistently with the stock market data, we use monthly average sales from 1978 to 2003. Note that we use sales instead of total circulation because the latter includes also subscriptions. We do not want to consider subscriptions to avoid capturing the behavior of professional investors, who typically get access to *Il Sole 24Ore* through annual subscriptions. (In Italy, subscriptions account for a very small part of newspapers sales anyway.)

It could reasonably be argued that a financial newspaper is not the only way to acquire information on the stock market. Other financial publications, national newspapers, and internet services could to some extent be substitute channels of information acquisition with respect to *Il Sole 24Ore*. However, there are good reasons to think that these alternative information sources are

[14] The source of these data is the ADS (Accertamenti Diffusione Stampa) dataset, which is publicly available (on paper) since 1976. ADS collects and certifies the publishers' declarations on the number of copies sold and printed and on the number of subscriptions.

[15] There are also two minor publications whose circulation is negligible (see figures quoted in two decisions of the Italian competition agency, Provv. n. 3336 *Class Editori / Il Sole 24 Ore* (19/10/1995) and Provv. n. 4822 *Italia Oggi Editori / Il Sole 24 Ore* (27/3/1997)).

[16] *Italia Oggi* and *MF* were founded in 1986 and 1989 respectively, so they were not present on the market in the first part of our sample period. Since *MF* is not present on the ADS dataset (because they chose not to report their data to this agency), the information available on this publication amounts to the figures provided by the Italian antitrust competition agency in two decisions (Provv. n. 3336 *Class Editori / Il Sole 24Ore* (19/10/1995) and Provv. n. 4822 *Italia Oggi Editori / Il Sole 24Ore* (27/3/1997)), according to which *MF* had a share of 2-3% in 1993 and in 1995. As to *Italia Oggi*, its market share was on average 6-7% over the whole period. Therefore *Il Sole 24Ore* had a market share of around 90% or higher in the period under consideration.

not crucial for the purposes of our analysis. First of all, national newspapers of general information do not seem to be close substitutes for *Il Sole 24Ore*, because of their lower coverage and level of in-depth analyses on financial information.[17]

With the advent of internet, online news have become an alternative channel with respect to written publications. However, since our data go back to 1978, internet was not even existing or widespread for most of the period under consideration. The development of internet access is a relatively recent phenomenon in Italy, and internet usage was fairly limited even until five years ago.[18,19] In the next section, we also perform the econometric analysis on a restricted sample where we have deleted the last three years of observations, in order to eliminate the period where internet started to be an alternative source of financial information.

Another source of information on financial markets is provided by private agencies (e.g. Reuters or Bloomberg) which charge a subscription fee for their service. Therefore they are generally targeted to professional investors who need to have constant and detailed information about instantaneous variations in the stock prices. The target readership of a financial newspaper like *Il Sole 24Ore* is instead mainly composed of individual investors who want to find more detailed information about the financial assets they hold, obtain relevant news, and read experts' analyses. Since individual, non-professional investors are more likely to exhibit a 'cognitive dissonant' behavior, we focus therefore on *Il Sole 24Ore*, which is probably the main source of information for this type of investors (and - as said above - we do not consider subscriptions, to focus on non-professional investors' decisions).

[17]The view that the market for daily financial newspapers can be defined as a separate market is taken also by the Italian competition agency in the two above-mentioned decisions.

[18]The share of internet users over the total population was around 1% in 1997, and started to increase sharply only after the year 2000, when it reached 22% of the population (source: Computer Industry Almanac). However, high-speed connections are not widespread even nowadays, and prices for internet access still limit the frequency of usage.

[19]Furthermore, even if individual investors can find free information about the general performance of the stock market in many websites, there are several reasons why printed media still provide a valuable service (see Dyck and Zingales, 2003). They include credibility of information as well as sound and in-depth analyses of the determinants and prospects of stock exchange performance.

In order to get a first glance at the data, in Figure 1 below we plot the series of sales of *Il Sole 24Ore* against the series of the Datastream index. In order to interpret the increasing trend of the stock market index, it should be taken into account that the Italian stock exchange has grown considerably in the 1980s and even more in the 1990s. Mutual funds were introduced in 1983, but it is only from the second half of the 1980s that they started to become a widely held financial instrument. Household participation in equity markets increased from 26.43% in 1985 to 38.19% in 1995 and to 48.24% in 1998 (source: Pelizzon and Weber, 2004, based on information from a Bank of Italy SHIW survey).

A first look at the raw data series seems to anticipate the results that we illustrate in the next section: this graph suggests that there is some kind of relation between the sales of *Il Sole 24Ore* and the stock market index. The investigation of the existence and the direction of this relationship is the object of the next section.

Figure 1: Sales of *Il Sole 24Ore* and stock market index (Note: the scale for Sole is on the left axis and the scale for the index is on the right axis).

3 Econometric Methodology and Results

3.1 Main findings

As explained earlier, the focus of interest is the relation between the *Sole 24Ore* sales and the stock price index. First, we perform a cointegration analysis. This is important because if two variables are cointegrated then they are trending together (are driven by the same stochastic trend) and hence there is a strong relation between them. Indeed, it can be shown that *two cointegrated variables must be Granger-causally related at least in one direction*. Then, we carry out Granger-causality tests based on VAR processes, in order to investigate the *direction* of causality (Granger, 1969).[20]

Since our methods depend on the unit root or integration and cointegration properties of the variables, we first perform unit root tests. We find support for classifying both variables as integrated of order one ($I(1)$), suggesting that they have a stochastic trend (see Appendix for a discussion of the unit root properties of the variables).

Cointegration is explored with Johansen's likelihood ratio trace tests (Johansen, 1995). These tests check the number of cointegration relations in a VAR framework and use the full information in all the time series involved. Since only two $I(1)$ variables are considered, the only null hypothesis of interest is no cointegration relation between them. If that hypothesis is rejected, we conclude that there is a cointegration relation.[21]

The results are given in Table 1 and they suggest a cointegration relation between (log) *Sole 24Ore* sales ("log Sole") and the (log) stock market price index ("log index"). We are using different lag orders in the tests because the test results are known to be sensitive to the number of lags included in

[20] Two time series variables x and y are Granger-causally related if one variable contains information for improving the forecasts of the other variable. More precisely, x is Granger-causal for y if the forecasts of y can be improved by taking into account past and present information in x.

[21] The tests are based on a model where no deterministic linear trend is allowed in the cointegration relation, that is, the linear trend (if it exists) is orthogonal to the cointegration relation. Notice that the variables may still have a linear trend individually.

Table 1: Johansen Trace Tests for one Cointegration Relation in the log Sole/log index System

variables (sample period)	deterministic term	VAR lag order	value of test stat.	p-value
log Sole/log index	orth. trend, sd	2	15.27	0.05
(1978M1-2003M12, $T = 312$)	orth. trend, sd	5	16.36	0.04
	orth. trend, sd	6	14.09	0.08

Note: sd - seasonal dummies.

the underlying VAR model.[22] For all three lag orders reported in Table 1, a cointegration rank of zero and, hence, the lack of cointegration can be rejected at least at the 10% level for the log Sole/log index system. *Therefore there is evidence of a cointegration relation between the Sole sales and the stock market index, which is consistent with the "cognitive dissonance" hypothesis.*

Since cointegration is a sufficient but not a necessary condition for Granger-causality, we next investigate the direction of causality between the variables using Wald tests for Granger-causality.[23] Test results based on VAR models with different lag orders and an intercept term as well as seasonal dummy variables as deterministic terms are presented in Table 2.[24,25] *These test results present a strong case in favor of a Granger-causal relation from the log stock market index to the log Sole sales.* More precisely, the null hypothesis of no Granger-causality from log index to log Sole (log index $\not\rightarrow$ log Sole) is clearly rejected at common significance levels. We have also checked that the existence of the inverse causality relationship is not confirmed by the data (see Appendix).

In Table 2 we also present results of tests for what is usually called instantaneous causality in the literature, and may be viewed as a measure for the in-

[22] Our choice of the lag orders is based on the three standard model selection criteria AIC, HQ and SC (e.g., Lütkepohl, 2005, Section 4.3).

[23] If cointegrated variables are involved, standard Wald tests (F-tests) for Granger-causality are problematic in general (Toda and Phillips, 1993). In the present situation it can be shown, however, that they have their usual asymptotic χ^2- or approximate F-distributions. The log Sole/log index system is one with two $I(1)$ variables and one cointegration relation for which the Wald statistic for Granger-causality has standard properties (see Toda and Phillips, 1993 and Lütkepohl and Reimers, 1992).

[24] The lag orders are again chosen by the three most common criteria in applied work (AIC, HQ, SC).

[25] Here the F-version of the Wald test is applied which corrects at least partly for the undesirable small sample properties of the usual χ^2-version (Lütkepohl, 2005, Section 3.6).

Table 2: Causality Tests Based on VAR Models in Levels of log Sole/log index with p Lags, Sample Period 1978M1-2003M12

p	H_0	Granger causality	Instantaneous causality
6	log index $\nrightarrow$ log Sole	2.48 (0.02)	3.37 (0.07)
5	log index $\nrightarrow$ log Sole	3.13 (0.01)	3.35 (0.07)
2	log index $\nrightarrow$ log Sole	3.87 (0.02)	2.34 (0.13)

Note: p-values in parentheses.

Table 3: Causality Tests Based on VAR Models in Levels of log Sole/log index with p Lags, Sample Period 1978M1-2000M12

p	H_0	Granger causality
6	log index $\nrightarrow$ log Sole	3.19 (0.00)
5	log index $\nrightarrow$ log Sole	3.88 (0.00)
2	log index $\nrightarrow$ log Sole	6.94 (0.00)

Note: p-values in parentheses.

stantaneous relation between the two variables when all intertemporal relations have been accounted for. The table shows that there is a (weak) instantaneous relation between the variables. Thus, the results of the instantaneous causality tests are consistent with the previous conclusion in favor of the cognitive dissonance hypothesis.

To check the robustness of our results, we have also deleted the data associated with the last three years of our sample and we have repeated the tests. By deleting the data after the year 2000 eliminates much of the long lasting downward movement in the stock market which some may regard as an unusual period. Moreover, the past few years also coincide with the advent of internet websites as an alternative source of financial information, which may weaken the relationship between stock market variables and *Sole 24Ore* sales. The results for the restricted sample are shown in Table 3. With respect to the results for the full sample, they show even stronger evidence of a causality relationship between the stock market index and *Sole 24Ore* sales.

We also performed an impulse response analysis, which confirms the positive relationship between the two series: when the stock market index goes down

Table 4: Johansen Trace Tests for one Cointegration Relation in log Sole/log volume and log index/log volume Systems

variables (sample period)	deterministic term	VAR lag order	value of test stat.	p-value
log Sole/log volume	orth. trend, sd	2	17.15	0.03
(1986M7-2003M12, $T = 210$)	orth. trend, sd	5	9.75	0.31
log index/log volume (1986M7-2003M12, $T = 210$)	orth. trend, sd	2	10.33	0.26

Note: sd - seasonal dummies.

the *Sole 24Ore* sales will also go down.

3.2 Alternative explanations and robustness checks

An alternative motivation for searching information on financial markets follows a more "standard" explanation, according to which investors may buy financial newspapers in order to improve their trading in the market. According to this view, one should expect to find that the financial newspaper sales increase with the volumes traded in the market.

However, the results reported in Table 4 do not suggest the existence of cointegration between newspaper sales ("log Sole") and traded volumes ("log volume"). The test results for these two variables depend strongly on the lag order chosen and the null hypothesis of no cointegration is not rejected at common significance levels for lag order 5.

Note also an important implication of a cointegration relation between the two pairs of variables. If both the stock price index and the volume were cointegrated with the newspaper sales, then they must necessarily also be cointegrated with each other (e.g., Lütkepohl, 2005, Section 6.3). Therefore cointegration between log index and log volume is also checked in Table 4 where it is seen that the rank zero hypothesis cannot be rejected. This implies that there exists no cointegration relation between *Sole 24Ore* sales and the stock market volume.

We next investigate the direction of Granger-causality between the stock market volume and *Sole 24Ore* sales. The system (log Sole/log volume) consists of two $I(1)$ variables which are not cointegrated. Hence, a stationary VAR in

Table 5: Causality Tests Based on VAR Models in First Differences of log Sole/log volume with p Lags

p	H_0	Granger causality Sample period 1978M1-2003M12	Granger causality Sample period 1978M1-2000M12
1	Δ log volume $\nrightarrow$ Δ log Sole	0.04 (0.85)	0.19 (0.66)
4	Δ log volume $\nrightarrow$ Δ log Sole	1.84 (0.12)	0.45 (0.77)

Note: p-values in parentheses.

first differences may be considered and used as the basis for Granger-causality tests.[26] The test results presented in Table 5 show that in both sample periods there is no evidence for a Granger-causal relation from the stock market volume to *Sole 24Ore* sales, since the null hypothesis of no Granger-causality cannot be rejected at usual significance levels.[27] (The Appendix shows that there is no causality link going from *Sole 24Ore* sales to stock market volumes, either.)

As we have said in the Introduction, some neoclassical models predict that the greater the magnitude of the noise in the stock market index, the higher the incentive to acquire information. However, we show that there is no evidence of a causal relation between stock market volatility and newspaper sales.

In order to explore this relationship, we have performed Granger-causality tests for a series of squared, mean-adjusted stock index returns and log Sole. More precisely, the new variable is $r_t^2 = (\Delta \log \text{index}_t - \hat{\mu})^2$, where $\hat{\mu} = T^{-1} \sum_{t=1}^{T} \Delta \log \text{index}_t$. This series may be viewed as a measure of market volatility. The series does not have a unit root and one could argue that checking Granger-causality between r_t^2 and the changes in log Sole is preferable to using log Sole in levels. Therefore both types of tests are reported in Table 6. Clearly, there is no evidence of a causal relation from market volatility to the newspaper sales.

Therefore our analysis does not lend support to these two particular "rational" explanations of information acquisition in financial markets.

We have also performed a series of robustness checks which are reported

[26] In stationary VARs the usual Wald tests are known to have standard asymptotic properties, however (Lütkepohl, 2005, Section 3.6).

[27] Results of the instantaneous causality tests are in line with those of Granger-causality tests.

Table 6: Causality Tests Based on VAR Models for Market Volatility and log Sole, Sample Period 1978M1-2003M12

p	H_0	Granger causality
9	$r^2 \nrightarrow$ log Sole	1.19 (0.30)
1	$r^2 \nrightarrow$ log Sole	0.00 (1.00)
9	$r^2 \nrightarrow \Delta$ log Sole	0.91 (0.52)
4	$r^2 \nrightarrow \Delta$ log Sole	0.63 (0.64)
1	$r^2 \nrightarrow \Delta$ log Sole	0.02 (0.90)

Note: p-values in parentheses.

in the Appendix. In particular, we find no or at best very little evidence of the existence of reverse causality relations between the variables. We have also performed lag augmentation tests à la Dolado and Lütkepohl (1996), whose results confirm our findings. Finally, we show results of causality tests for the levels VAR models for the log Sole/log index system with a deterministic linear trend term in addition to seasonal dummy variables. Again they confirm the results obtained with a constant and seasonal dummies only.

One could argue that our finding of a relationship between the stock market index and the financial newspaper sales may be due to an omitted factor rather than to the cognitive dissonance explanation. For example, one could think that when the stock market goes up, people feel richer and increase their consumption, including the purchase of *any* type of newspaper. If this is the case, we should therefore find a relationship between the stock market index and the sales of non-financial newspapers. We therefore performed a cointegration analysis between the stock price index and the sales of Italian sports newspapers (which are probably the class of newspapers more distant from the financial ones), whose details are shown in the Appendix. Our analysis shows that there is no evidence of cointegration between these variables, nor evidence of a causal relation between them. This supports the idea that there is no spurious correlation between the newspaper sales and the stock market index, which somewhat lends more support to the behavioral explanation.

Another variable which might be used to check whether the relation of interest is spurious is an index of economic activity. In the Appendix, we show results of a causality analysis between Italian industrial production (used as a proxy for economic activity) and *Sole 24Ore* sales. The two series are not cointegrated nor there is trace of a causal relation (neither Granger nor instantaneous) between them.

4 Results on the UK data

In order to provide additional evidence of cognitive dissonance, we replicate the analysis using UK data. We therefore consider the relationship between the series of monthly average circulation for the Financial Times (UK)[28] and the series of the (monthly average) Datastream index for the UK stock market.

The plot of the two series is already suggesting that the two series move together. It is quite striking how much the FT sales trace the UK index. Notice however that the plots cover the full sample period starting in 1978, including the years with missing values. In the following analysis only data from 1985M1 - 2007M5 are used. In other words we have only used the period without missing observations.

[28] In the following the Financial Times circulation is abbreviated as FT.

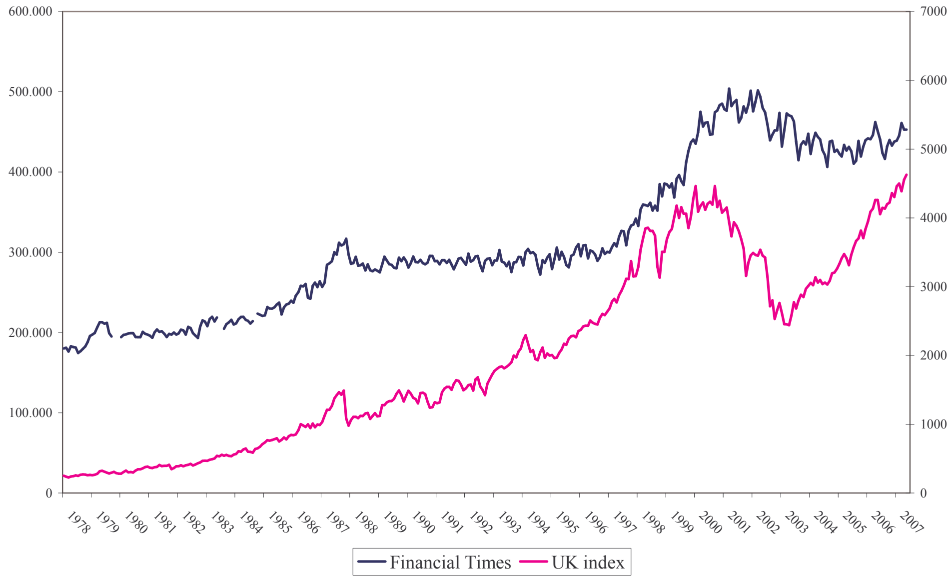

Figure 2: Financial Times circulation and UK stock market index (Note: the scale for the FT is on the left axis and the scale for the index is on the right axis).

Results of cointegration tests are given in Table 7. In this case the Johansen trace tests do not show clear evidence of a cointegration relation. Therefore we have also applied the potentially more powerful Saikkonen/Lütkepohl tests[29] (S&L in the table) and they clearly reject the null hypothesis of no cointegration at least at the 10% level, depending on the VAR order used. In any case, these results can be interpreted as moderately strong evidence in favor of cointegration between the log FT and log UK index series.

As expected on the basis of the plots, strong evidence for a causal relation from log UK index to log FT is found based on the causality tests presented in Tables 8 and 9. As in the Italian case, there is again no evidence of a reverse Granger causality relation (see Appendix). Other robustness checks, reported in the Appendix, confirm the strong causal relation from the log UK index to

[29] A power comparison between the Johansen tests and the Saikkonen/Lütkepohl tests is presented in Saikkonen and Lütkepohl (2000).

Table 7: Trace Tests for one Cointegration Relation in the log FT/log UK index System, Sample Period 1985M1-2007M5, $T = 269$

test	deterministic term	VAR lag order	value of test stat.	p-value
Johansen	orth. trend, *sd*	2	12.04	0.16
	orth. trend, *sd*	3	12.73	0.13
S&L	orth. trend, *sd*	2	9.55	0.06
	orth. trend, *sd*	3	10.05	0.05

Note: *sd* - seasonal dummies.
Order selection by SC, HQ and AIC criteria with maximal lag 24.

the log Financial Times sales. Overall, the results for the UK market are fully in line with those for the Italian market.

Table 8: Causality Tests Based on VAR Models in Levels of log FT/log UK index with p Lags, Sample Period 1985M1-2007M5, $T = 269$

p	H_0	Granger causality	Instantaneous causality
2	log UK index $\nrightarrow$ log FT	8.56 (0.00)	4.65 (0.03)

Note: p-values in parentheses.

Table 9: Causality Tests Based on VAR Models in Levels log FT/log UK index with p Lags, Sample Period 1985M1-2000M12

p	H_0	Granger causality
2	log UK index $\nrightarrow$ log FT	6.00 (0.00)

Note: p-values in parentheses.

5 Conclusions

In this paper we have tested a prediction about (non-professional) agents' attitude toward information acquisition according to which individuals may display a 'cognitive dissonant' behavior by refusing to acquire available information which might contrast with their maintained beliefs.

We analyze this issue by looking at the relationship between the stock market and the demand for financial information in Italy. If agents' behavior is driven by psychological considerations, we should expect them to purchase the newspaper (and therefore to acquire information) when the stock market performs well (because they are more likely to find that their shares are doing well) and not to purchase it when the stock market is in a negative phase. Therefore according to this hypothesis we should find that an increase in the stock market price increases the sales of the financial newspaper.

A cointegration analysis on the series of the Italian stock exchange price index and of the Italian main financial newspaper's circulation shows that the two series are indeed cointegrated. The Granger-causality relation between the two variables has the expected direction: we find evidence of a causality relation that goes from the stock market index to the sales of the newspaper, which is consistent with the 'cognitive dissonance' hypothesis.

According to 'standard' economic theory, agents should be eager to search for (relevant) information in order to improve their decision-making. If agents behave according to standard economic theory, they should acquire information on the financial market for transactional reasons (for instance, with a view to reoptimise their portfolio allocation). However, we do not find evidence of a cointegration relation between the volumes of transaction on the stock market and the demand for information (that we measure in terms of sales of the main Italian financial newspaper), nor evidence of a Granger-causal relation.

We also test another prediction of the rational model according to which the demand for financial information increases with stock price noise, but we do not find support for this hypothesis. Therefore our analysis does not lend support to these two 'rational' explanations of information acquisition in financial markets.

References

[1] Akerlof, G.A. and W.T. Dickens (1982), "The Economic Consequences of Cognitive Dissonance", *American Economic Review*, 72(3): 307-319.

[2] Benabou, R. and J. Tirole (2002), "Self-Confidence and Personal Motivation", *Quarterly Journal of Economics*, 470: 871-915.

[3] Blackwell, D. (1951), "Comparison of Experiments", in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, Berkeley, CA, 93-102.

[4] Blackwell, D. (1953), "Equivalent Comparison of Experiments", *Annals of Mathematical Statistics*, 24: 265-272.

[5] Brocas, I. and J.D. Carrillo (eds.) (2003), *The Psychology of Economic Decisions*, Vol.1, Oxford University Press.

[6] Brocas, I. and J.D. Carrillo (eds.) (2004), *The Psychology of Economic Decisions*, Vol.2, Oxford University Press.

[7] Camerer, C.F., G. Loewenstein, and M. Rabin (eds.) (2004), *Advances in Behavioral Economics*, Princeton University Press.

[8] Carrillo, J.D. and T. Mariotti (2000), "Strategic Ignorance as a Self-Disciplining Device", *Review of Economic Studies*, 67: 529-544.

[9] Dolado, J. J. and H. Lütkepohl (1996), "Making Wald Tests Work for Cointegrated VAR Systems", *Econometric Reviews,* 15: 369–386.

[10] Dyck, A. and L. Zingales (2003), "The Media and Asset Prices", mimeo.

[11] Eliaz, K. and R. Spiegler (2003), "Anticipatory Feelings and Choices of Information Sources", mimeo, NYU and Tel Aviv University.

[12] Fama, E.F. and K.R. French (1988), "Permanent and Temporary Components of Stock Prices", *Journal of Political Economy*, 96: 246-273.

[13] Festinger, L. (1957), *A theory of cognitive dissonance*, Evanston, IL: Row, Peterson.

[14] Granger, C. W. J. (1969), "Investigating Causal Relations by Econometric Models and Cross-spectral Methods", *Econometrica*, 37: 424–438.

[15] Grossman, S.J. and J.E Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets", *American Economic Review*, 70(3): 393-408.

[16] Guiso, L. and T. Jappelli (2005), "Rational or Behavioral Households? Investment in Information and Portfolio Performance", mimeo.

[17] Harmon-Jones, E. and J. Mills (eds.) (1999), *Cognitive Dissonance: Progress on a Pivotal Theory in Social Psychology*, American Psychological Association.

[18] Hirshleifer, J. (1971), "The Private and Social Value of Information and the Reward to Inventive Activity", *American Economic Review,* 61(4): 561-574.

[19] Johansen, S. (1995), *Likelihood-based Inference in Cointegrated Vector Autoregressive Models*, Oxford University Press, Oxford.

[20] Karlsson, N., G. Loewenstein, and D. Seppi (2005), "The 'Ostrich Effect': Selective Exposure to Information About Investments", downloadable from www.ssrn.com.

[21] Köszegi B. (2001), "Ego Utility and Information Acquisition", UC Berkeley, mimeo.

[22] Lütkepohl, H. (2005), *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin.

[23] Lütkepohl, H. and M. Krätzig (eds) (2004), *Applied Time Series Econometrics*, Cambridge University Press, Cambridge.

[24] Lütkepohl, H. and H.-E. Reimers (1992), "Granger-causality in Cointegrated VAR Processes: The Case of the Term Structure", *Economics Letters,* 40: 263–268.

[25] Odean, T. (1998), "Are Investors Reluctant to Realize their Losses?", *The Journal of Finance*, 53: 1775-1798.

[26] Osborne, M.J. (2004), *An Introduction to Game Theory*, Oxford University Press, New York.

[27] Pelizzon, L. and G. Weber (2004), "Are Household Portfolio Efficient? An Analysis Conditional on Housing", forthcoming on the *Journal of Financial Quantitative Analysis.*

[28] Peress, J. (2004), "Wealth, Information Acquisition, and Portfolio Choice", *Review of Financial Studies*, 17(3): 879-914.

[29] Rabin, M. (1998), "Psychology and Economics", *Journal of Economic Literature*, 36: 11-46.

[30] Rabin, M. (2002), "A Perspective on Psychology and Economics", *European Economic Review*, 46: 657-685.

[31] Rabin, M. and J.L. Schrag (1999), "First Impression Matters: A Model of Confirmatory Bias", *Quarterly Journal of Economics*, 114: 37-82.

[32] Saikkonen, P. and H. Lütkepohl (2000), "Testing for the Cointegrating Rank of a VAR Process with an Intercept", *Econometric Theory*, 16: 373-406.

[33] Toda, H. Y. and P. C. B. Phillips (1993), "Vector autoregressions and causality", *Econometrica,* 61: 1367–1393.

[34] Yariv, L. (2005), "I'll See It When I Believe It - A Simple Model of Cognitive Consistency", California Institute of Technology, mimeo.

Appendix

Unit root properties of the variables

We have used augmented Dickey-Fuller (ADF) and KPSS tests (see Lütkepohl and Krätzig, 2004, Section 2.7) to formally investigate the unit root properties of the variables. Some results are presented in Table 10 below.[30] They suggest that all three series may indeed be classified as $I(1)$. ADF tests of the levels of all series cannot reject the null hypothesis of a unit root while KPSS tests of stationarity clearly reject. Furthermore, ADF tests of the first differences clearly reject a unit root and thereby confirm that higher order integration can be excluded.

Table 10: Unit Root Tests for log Sole, log index and log volume

variable	sample period	test	deterministic terms	lag order	value of test stat.
log Sole	1978M1-2003M12	ADF	c, sd, t	4	-0.59
log Sole	1978M1-2003M12	KPSS	c, t	12	0.55^{**}
Δlog Sole	1978M2-2003M12	ADF	c, sd	2	-12.0^{**}
log index	1978M1-2003M12	ADF	c, t	3	-2.23
log index	1978M1-2003M12	ADF	c	3	-2.38
log index	1978M1-2003M12	KPSS	c, t	12	0.33^{**}
Δlog index	1978M1-2003M12	ADF	c	12	-3.89^{**}
log volume	1986M7-2003M12	ADF	c, t	3	-2.44
log volume	1986M7-2003M12	ADF	c	1	-1.28
log volume	1986M7-2003M12	KPSS	c, t	12	0.23^{**}
Δlog volume	1986M8-2003M12	ADF	c	3	-19.1^{**}

Notes: c - constant, t - linear time trend, sd - seasonal dummies.
Asterisks * and ** indicate that the null hypothesis of a unit root can be rejected at the 5% and 1% levels, respectively.
Order selection for ADF test by HQ criterion with maximal lag 24.

[30] All computations were performed with the software JMulTi (Lütkepohl and Krätzig, 2004).

Table 11: Causality Tests Based on VAR Models in Levels of log Sole/log index and First Differences of log Sole/log volume with p Lags

p	H_0	Granger causality Sample period 1978M1-2003M12	Granger causality Sample period 1978M1-2000M12
6	log Sole $\not\rightarrow$ log index	2.78 (0.01)	1.22 (0.30)
5	log Sole $\not\rightarrow$ log index	1.53 (0.18)	0.72 (0.61)
2	log Sole $\not\rightarrow$ log index	2.00 (0.14)	1.03 (0.36)
1	Δ log Sole $\not\rightarrow$ Δ log volume	0.50 (0.48)	0.67 (0.41)
4	Δ log Sole $\not\rightarrow$ Δ log volume	0.60 (0.66)	0.90 (0.46)

Note: p-values in parentheses.

Additional causality tests

We have also checked that the existence of the inverse causality relationships is not confirmed by the data. Table 11 shows that there is no evidence of a causality direction going from log Sole to log index (at least for lag orders $p = 2$ and 5) and from log Sole to log volumes.

In order to check the robustness of our results, we have also used lag augmentation tests à la Dolado and Lütkepohl (1996) which are generally valid for integrated and cointegrated systems. The results for the full sample period are presented in Table 12 and they confirm our more refined findings. The advantage of these tests is, however, that they would also be asymptotically valid if the cointegration properties of the system were misspecified.

In Table 13 we also show results of causality tests for the levels VAR models for the log Sole/log index system which include a deterministic linear trend term in addition to seasonal dummy variables. Again they confirm our results obtained with a constant and seasonal dummies only. Given the way the time series look like, deterministic trend terms do not make much sense for models in first differences. Therefore we do not report such results.

Causality analysis for sports newspaper sales and stock market index

In analyzing the causal relation between the index and the sports newspaper sales we have used only data from 1987 to 2003 because there are data problems

Table 12: Lag Augmented Causality Tests for log Sole/log index and log Sole/log volume Systems, Sample Periods 1978M1-2003M12 and 1986M7-2003M12, respectively

p	H_0	Granger causality	instantaneous relation
2	log index $\not\rightarrow$ log Sole	4.377 (0.01)	1.972 (0.16)
6	log index $\not\rightarrow$ log Sole	2.548 (0.02)	3.444 (0.06)
2	log volume $\not\rightarrow$ log Sole	0.013 (0.99)	1.954 (0.16)
6	log volume $\not\rightarrow$ log Sole	1.475 (0.19)	3.322 (0.07)

Note: p-values in parentheses.

Table 13: Causality Tests Based on VAR Models in Levels of log Sole/log index with Linear Trends and p Lags, Sample Period 1978M1-2003M12

p	H_0	Granger causality	instantaneous relation
6	log index $\not\rightarrow$ log Sole	2.41 (0.03)	3.44 (0.06)
5	log index $\not\rightarrow$ log Sole	3.00 (0.01)	3.38 (0.07)
2	log index $\not\rightarrow$ log Sole	3.86 (0.02)	2.34 (0.13)

Note: p-values in parentheses.

for the sports newspaper sales before 1987 (referred to as “Sports” and transformed in logs). We consider the aggregated sales of the three Italian sports newspapers which are available in our dataset (*Corriere dello Sport, Gazzetta dello Sport, Tuttosport*).

Unit root tests support a single nonseasonal unit root in the series which will therefore be treated as $I(1)$. Thus it makes sense to consider cointegration between log index and log Sports. The results of Johansen tests are given in Table 14. No cointegration is found, that is, the null hypothesis of no cointegration relation cannot be rejected.

Because cointegration is just a sufficient condition but not a necessary one for Granger-causality, we have also performed causality tests based on a VAR model for the first differences of the two series. The results are presented in Table 15. Clearly, there is no evidence for a causal relation in either direction.

Table 14: Johansen Trace Tests for one Cointegration Relation between log Sports and log index (H_0 : no cointegration)

variables (sample period)	deterministic terms	VAR lag order	value of test stat.	p-value
log Sports/log index	orth. trend, sd	2	6.95	0.59
(1987M1-2003M12, $T = 204$)	orth. trend, sd	1	9.50	0.33

Notes: sd - seasonal dummies.

These results support the idea that the relation between log index and log Sole is not just spurious.

Table 15: Causality Tests Based on a VAR(1) Model for log index and log Sports, Sample Period 1987M2-2003M12, $T = 203$

p	H_0	Granger causality
1	Δ log index $\not\rightarrow$ Δ log Sports	0.27 (0.60)
1	Δ log Sports $\not\rightarrow$ Δ log index	0.26 (0.61)

Note: p-values in parentheses.

Causality analysis for economic activity and newspaper sales

Causality tests between Italian industrial production ("IP") and *Sole 24Ore* sales are performed on a VAR(12) model with constant and seasonal dummies for the first differences of the logs. The two log series are not cointegrated. Table 16 shows that there is no trace of a causal relation (neither Granger nor instantaneous) in either direction between the two series.

Additional tests for the UK data

The unit root tests suggest that both series may indeed be classified as $I(1)$. ADF tests of the levels of both series cannot reject the null hypothesis of a unit root while KPSS tests of stationarity clearly reject. Furthermore, ADF tests of the first differences clearly reject a unit root and thereby confirm that higher order integration can be excluded.

Table 16: Causality Tests Based on VAR Models for Industrial Production and log Sole, Sample Period 1978M2-2003M12, $T = 311$

p	H_0	Granger causality	Instantaneous causality
12	Δ log IP $\nrightarrow$ Δ log Sole	0.57 (0.86)	1.08 (0.30)
12	Δ log Sole $\nrightarrow$ Δ log IP	0.82 (0.63)	1.08 (0.30)

Note: p-values in parentheses.

Table 17: Unit Root Tests for log FT and log UK index

variable	sample period	test	deterministic terms	lag order	value of test stat.
log FT	1985M1-2007M5	ADF	c, sd, t	2	-1.35
log FT	1985M1-2007M5	KPSS	c, t	12	0.19^{*}
Δlog FT	1985M2-2007M5	ADF	c, sd	1	-22.4^{**}
log UK index	1985M1-2007M5	ADF	c, t	1	-2.19
log UK index	1985M1-2007M5	ADF	c	1	-1.57
log UK index	1985M1-2007M5	KPSS	c, t	12	0.32^{**}
Δlog UK index	1985M2-2007M5	ADF	c	1	-15.7^{**}

Notes: c - constant, t - linear time trend, sd - seasonal dummies.
Asterisks * and ** indicate that the null hypothesis of a unit root can be rejected at the 5% and 1% levels, respectively.
Order selection for ADF test by HQ criterion with maximal lag 24.

Table 18 shows that there is no evidence of a reverse Granger causality relation between the two variables. Also the robustness checks in Tables 19 and 20 confirm the strong causal relation from the log UK index to the log Financial Times sales.

Table 18: Causality Tests Based on VAR Models in Levels of log FT/log UK index with p Lags, Sample Period 1985M1-2007M5

p	H_0	Granger causality	instantaneous relation
2	log FT $\not\rightarrow$ log UK index	0.63 (0.53)	4.65 (0.03)

Note: p-values in parentheses.

Table 19: Lag Augmented Causality Tests of log FT/log UK index, Sample Period 1985M1-2007M5

p	H_0	Granger causality	instantaneous relation
2	log UK index $\not\rightarrow$ log FT	5.85 (0.00)	4.84 (0.03)

Note: p-values in parentheses.

Table 20: Causality Tests Based on VAR Models in Levels of log FT/log UK index with Linear Trends and p Lags, Sample Period 1985M1-2007M5

p	H_0	Granger causality	instantaneous relation
2	log UK index $\not\rightarrow$ log FT	9.50 (0.00)	6.23 (0.01)

Note: p-values in parentheses.