ELSEVIER Journal of Financial Economics 53 (1999) 189–216

JOURNAL OF Financial ECONOMICS

www.elsevier.com/locate/econbase

How are stock prices affected by the location of trade?☆

Kenneth A. Froot[a,b,*], Emil M. Dabora[c]

[a]Graduate School of Business Administration, Harvard University, Boston, MA 02163, USA
[b]National Bureau of Economic Research, Cambridge, MA 02138, USA
[c]Morgan Stanley & Co., New York, NY 10036, USA

Received 19 July 1996; received in revised form 15 July 1998

Abstract

We examine pairs of large, 'Siamese twin' companies whose stocks are traded around the world but have different trading and ownership habitats. Twins pool their cash flows, so, with integrated markets, twin stocks should move together. However, the difference between the prices of twin stocks appears to be correlated with the markets on which they are traded most, i.e., a twin's relative price rises when the market on which it is traded relatively intensively rises. We examine several explanations of this phenomenon including: the discretionary use of dividend income by parent companies: differences in parent expenditures; voting rights; currency fluctuations; ex-dividend date timing issues; and tax-induced investor heterogeneity. Only the last hypothesis can explain some, but not all, of the empirical facts. We conjecture that: (a) country-specific sentiment shocks might affect share intensity, (b) investors are rational, but markets are segmented by frictions other than international transactions costs, such as agency problems.

JEL classification: G150

Keywords: International equities; Anomalies; Closed-end funds

---

* Corresponding author. Tel.: 617 495 6677; fax: 617 496 7357.

E-mail address: kfroot@hbs.edu (K.A. Froot)

☆We thank Richard Meyer, André Perold, Leonard Rosenthal, Rick Ruback, Jeremy Stein, an anonymous referee, and numerous practitioners for helpful conversations and commentary, Chris Allen and Philip Hamilton for help in obtaining data, and the Q-Group, the New York Stock Exchange, and the Division of Research at Harvard Business School for financial support. All errors and opinions are our own.

PII: S0304-405X(99)00020-3

1. Introduction

The classical finance paradigm predicts that an asset's price is unaffected by its location of trade. If international financial markets are perfectly integrated, then a given set of risky cash flows has the same value and risk characteristics when its trade is redistributed across markets and investors.

This paper provides a stark example in which the location of trade and ownership appears to influence prices. We show that the stock prices of three of the world's largest and most liquid multinational companies are strongly influenced by locational factors. Specifically, we test whether location matters by examining 'Siamese-twin' company stocks, or pairs of corporations whose charter fixes the division of current and future equity cash flows to each twin. The twins each have their own stock, with its own distinct trading habitat. We examine three examples of Siamese twins: Royal Dutch Petroleum and Shell Transport and Trading, PLC; Unilever N.V. and Unilever PLC; and SmithKline Beecham. At face value, twin charters imply that the twins' stock prices should move in lockstep, in a ratio given by the proportional division of cash flows. Surprisingly, the stock prices of twins do not behave in this manner. Rosenthal and Young (1990) show that the stock prices of Royal Dutch-Shell and Unilever N.V./PLC exhibit persistent and strikingly large deviations from the ratio of adjusted cash flows. To this, we add that the stock prices of SmithKline Beecham exhibit similar types of deviations.

The main contribution of this paper is to show that the relative price of twin stocks is highly correlated with the relative stock-market indexes of the countries where the twins' stocks are traded most actively. For example, when the U.S. market moves up relative to the U.K. market, the price of Royal Dutch (which trades relatively more in New York) tends to rise relative to the price of its twin Shell (which trades relatively more in London). Similarly, when the dollar appreciates against the pound, the price of Royal Dutch tends to increase relative to that of Shell. We consider a number of obvious potential explanations for this behavior, but find that none is able to fully explain it.

A similar sort of phenomenon occurs with closed-end country funds, which invest in emerging markets but are financed by issuing shares on developed-country markets. It is well known that the prices of these shares differ from the net asset values of the fund portfolios. In particular, it appears that closed-end fund share prices comove most strongly with the stock market on which they trade, while net asset values comove most strongly with their local stock markets.[1]

[1] Hardouvelis et al. (1995) chronicle the behavior of 35 country funds. They find that the funds trade, on average, at a discount and that fund discounts are sensitive to movements in the host country, U.S. and world stock markets. Similarly, Bodurtha et al. (1993) find that the movement of closed-end country funds prices on U.S. markets is correlated with the U.S. market, while the underlying share prices are correlated with the foreign markets on which they trade. These papers build on Lee et al. (1991), which argues that closed-end fund discounts reflect the sentiment on small stocks (see also Chen et al., 1993; Chopra et al., 1993).

We believe our Siamese-twin stocks provide a more clear-cut example of 'excess comovement' for several reasons. First, the twins we examine are among the largest and most liquid stocks in the world. By contrast, closed-end funds (and many of the stocks they hold) are relatively illiquid, so their prices are not as 'clean'. Second, our Siamese-twin stocks represent claims on exactly the same underlying cash flows. Closed-end shares, on the other hand, are claims not only to a portfolio of foreign stocks, but also to the dynamic trading strategy followed by fund managers. The differences between fund share prices and net asset values might be explained by the perceived value of this strategy. Third, arbitrage between closed-end fund shares and net assets is costly or even forbidden.[2] Indeed, closed-end funds profit by enabling investors to better internationalize their portfolios, so funds tend to open where investment barriers are relatively high. By contrast, the stocks of our twins can be arbitraged easily. They trade on major world stock exchanges, and the twins' stock can both be purchased locally by many investors. For example, a U.S. (Dutch) investor can buy Royal Dutch *and* Shell in New York (Amsterdam). As a consequence, the *additional* costs and informational advantages commonly associated with cross-border trading cannot be used to explain our results.[3]

What sources of international segmentation might explain our findings? One hypothesis, which we discuss below, is that of cross-border tax rules. Withholding taxes on dividends differ across countries and investor clienteles. In most instances, however, the withholding taxes for any given investor are the same for the stocks of any pair of twins. Thus, while helpful, tax-driven stories cannot fully account for our findings.

A second possible source of segmentation is country-specific noise. Suppose that a noise shock hitting, say, U.S. stocks, disproportionately affects the twin which trades relatively more in New York. In other words, stocks that trade more actively in the local market are more sensitive to local noise shocks and less sensitive to foreign noise shocks. This story has an interesting implication: the component of market movements explained by changes in twin's relative prices is likely to be noise. Twin price disparities, which are readily observable, may therefore be informative about market-wide noise shocks, which are not directly observable.

Finally, the comovement patterns we observe might result from institutional frictions involving informational and contractual inefficiencies. Principals must

[2] Pontiff (1993) shows that the size and persistence of closed-end fund discounts are cross-sectionally related to measures of arbitrage costs between the net asset values and the fund shares.

[3] This argument assumes that the law of one price holds around the world for each stock. Our data support this assumption, as each individual stock trades for approximately the same price in all markets at the same time.

control the agents who invest on their behalf. To do this, it might be optimal to narrowly define agents' discretionary authority or to write contracts that provide incentives for agents to limit discretion. As a result, equity fund managers may be restricted to invest in U.S. or international stocks, or they may be benchmarked against a widely accepted index, such as the S&P 500 (which includes Royal Dutch and Unilever N.V.) or the Financial Times Allshare index (which includes Shell and Unilever PLC), even if that index does not exhibit optimal risk/return characteristics. All else equal, these arrangements can create a bias toward certain stocks and away from others, but the arrangements could be optimal given the information and agency problems in investing.

The rest of this paper is organized as follows. Section 2 briefly describes the organizational structure of the twins. Section 3 presents our tests of comovement and cointegration of price twin differentials. Section 4 discusses the data. Section 5 presents our findings on comovement. Section 6 discusses several possible explanations for the results. Section 7 offers conclusions.

2. The relations between pairs of corporate twins

2.1. Royal Dutch Petroleum and Shell Transport and Trading, PLC

Royal Dutch and Shell are independently incorporated in the Netherlands and England, respectively. The structure has grown 'out of a 1907 alliance' between Royal Dutch and Shell Transport by which the two companies agreed to merge their interests on a 60 : 40 basis while remaining separate and distinct entities (Royal Dutch 20 F, 1994, p. 1). All sets of cash flows, adjusting for corporate tax considerations and control rights, are effectively split in the proportion of 60 : 40.[4] Information clarifying the linkages between the two parent companies is widely available. In addition to being explained at the beginning of each Annual Report, the connections are detailed in 20F submissions to the SEC and are the subject of an analyst/investor guide (Royal Dutch Shell, 1994). There is also considerable public information about the relative pricing of Royal Dutch and Shell, and 'switch' trades are known by traders as those which seek to take advantage of price disparities between Royal Dutch and Shell.

[4] Royal Dutch and Shell Transport shall share in the aggregate net assets and in the net aggregate dividends and interest received from Group companies in the proportion of 60 : 40. It is further arranged that the burden of all taxes in the nature of or corresponding to an income tax leveeable in respect of such dividends and interest shall fall in the same proportion (Royal Dutch 20F, 1993, pp. 1–2). See also Rosenthal and Young (1990).

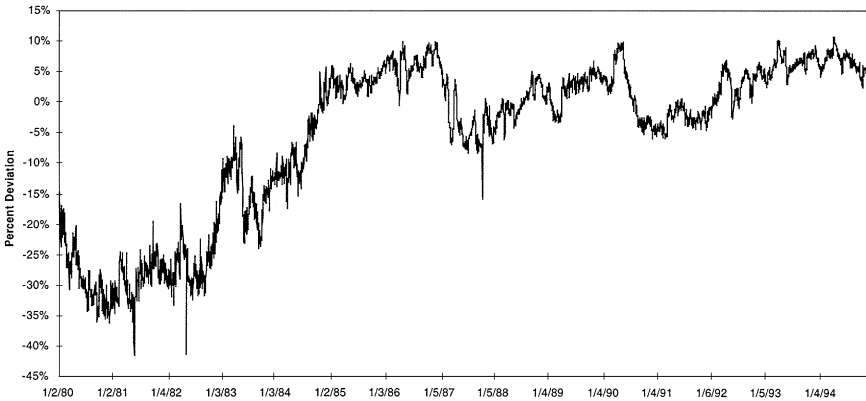

Fig. 1. Log deviations from Royal Dutch/Shell parity. Note: This figure shows on a percentage basis the deviations from theoretical parity of Royal Dutch and Shell shares and ADRs traded on the NYSE. Data are from the Center for Research in Security Pricing (CRSP).

Royal Dutch and Shell trade on nine exchanges in Europe and the U.S., but Royal Dutch trades primarily in the U.S. and the Netherlands (it is in the S&P 500 and virtually every index of Dutch shares) and Shell trades predominantly in the U.K. (it is in the Financial Times Allshare Index, or FTSE). Geographical ownership and trading information for Royal Dutch and Shell is shown in Table 1. Log deviations from the expected price ratio are graphed in Fig. 1.

2.2. Unilever N.V. and Unilever PLC

Unilever N.V. and Unilever PLC are independently incorporated in the Netherlands and England, respectively. In 1930, the two companies established an equalization agreement of cash flows. According to this agreement, the two companies act as a single group company and use the same board of directors. In the case of liquidation, all assets are to be pooled and divided evenly among shareholders. The intent of the agreement is to make the shares as similar as possible, as if all shareholders held shares of a single company. The Equalization Agreement states that distributions are 'made on the basis that the sum paid as dividends on every 1 pound nominal amount of PLC capital is equal . . . to the sum paid as dividends on every 12 fl. nominal amount of ordinary capital of N.V.' The PLC shares are listed as 5 pence per share, and the N.V. shares are listed at 4 fl per share. Thus earnings per share (expressed in a common currency) are equated by (1/5) PLC EPS = (12/4) N.V. EPS.[5]

[5] The 1993 Unilever N.V. 20F submission to the SEC (1993, p. 2) states: 'Since 1930 N.V. and PLC have operated as nearly as is practical as a single entity . . . they have agreed to cooperate in every way for the purpose of maintaining a common policy in every field of operations'. See also Rosenthal and Young (1990).

Unilever trades on eight exchanges in Europe and the U.S. N.V. trades mostly in the Netherlands, then in Switzerland and the U.S. (it is in the S&P 500). PLC trades predominantly in the U.K. (it is in the FTSE). Geographical ownership data are given in Table 1. Log deviations from the expected price ratio are graphed in Fig. 2.

2.3. SmithKline Beecham

SmithKline Beckman and Beecham Group merged to form SmithKline Beecham on July 26, 1989. The former holders of Beecham (a U.K. company) received class A ordinary shares while former holders of SmithKline Beckman (a U.S. corporation) received Equity Units (class E shares) comprised of 5 shares of SmithKline Beecham B ordinary shares and one preferred share of SmithKline Beecham Corporation. The equity units receive their dividends from SB Corp., a wholly owned American subsidiary. The dividends are equalized, so that one class E share provides the same dividend flow as one class A share.[6]

Geographic ownership data are unavailable, so Table 1 lists trading as a percentage of yearly trading volume. A shares are traded predominantly in the U.K., while H (the ADR on A shares) and E shares are traded in the U.S. Log deviations from parity are graphed in Fig. 3.

3. Empirical hypotheses and tests

Our null hypothesis is that relative twin prices should be uncorrelated with everything. Our alternative hypothesis is that markets are segmented, so that relative market shocks explain movements in the price differential. Specifically, we hypothesize that stocks that are most intensively traded on a given market will comove excessively with that market's return and currency.

To measure the relative comovement of twin prices, we regress the twins' log return differential on U.S., U.K., and Dutch market index log returns plus the relevant log currency changes:

$$r_{\mathrm{A-B},t} = \alpha + \sum_{i=-1}^{1} \beta_i S\&P_{t+i} + \sum_{j=-1}^{1} \delta_j FTSE_{t+j} + \sum_{k=-1}^{1} \lambda_k DI_{t+k} + \sum_{l=-1}^{1} \gamma_l gl/\$_{t+l} + \sum_{m=-1}^{1} \upsilon_m gl/£_{t+m} + \varepsilon_t, \quad (1)$$

[6] Dividends on Equity Units, which are paid by SmithKline Beecham Corporation ('SB Corp.'), are equivalent to the dividends on the A shares of the Company together with the related tax credit, and include the cumulative preference dividends on the Participating Preferred Shares of SB Corp. up to the date of payment…', (SmithKline Beecham Annual Report and Accounts, 1993).

Table 1
Distribution of share ownership and trading volume across markets

Panel A: Ownership (average 1980–1992)

Company	Percent owned in		
	U.S.	U.K.	Netherlands
Royal Dutch	33%	4%	34%
Shell	3%	96%	< 1%
Unilever N.V.	16%	10%	46%
Unilever PLC	< 1%	99%	< 1%

Panel B: Trading volume (average 1991–1995)

Company	Percent of average daily volume traded in		
	U.S.	U.K.	Netherlands
Royal Dutch	70%	NA	30%
Shell (ADR)	32%	68%	NA
SmithKline	83%	17%	NA

Sources: Royal Dutch and Shell 20-F statements, 1980–1992; Unilever N.V., 20-F, 1980–1983; a booklet published by Unilever N.V. entitled 'Charts 1984–1994'; Trading volume data are from the NYSE and London Stock Exchange.

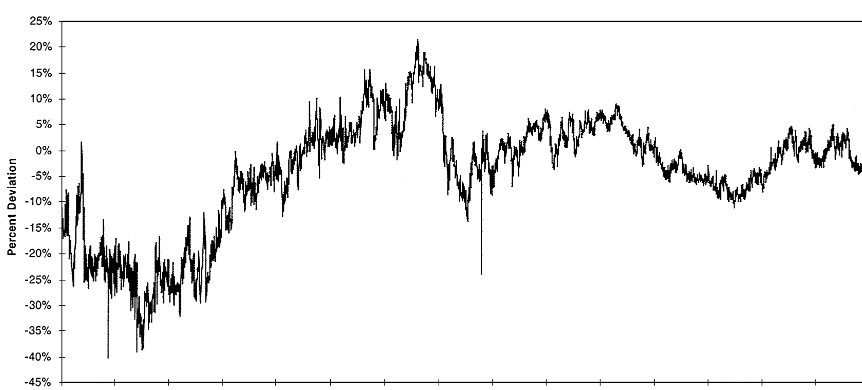

Fig. 2. Log deviations from Unilever N.V./Unilever PLC parity. Note: This figure shows on a percentage basis the deviations from theoretical parity of Unilever N.V. and PLC shares and ADRs traded on the NYSE. Data are from the Center for Research in Security Pricing (CRSP).

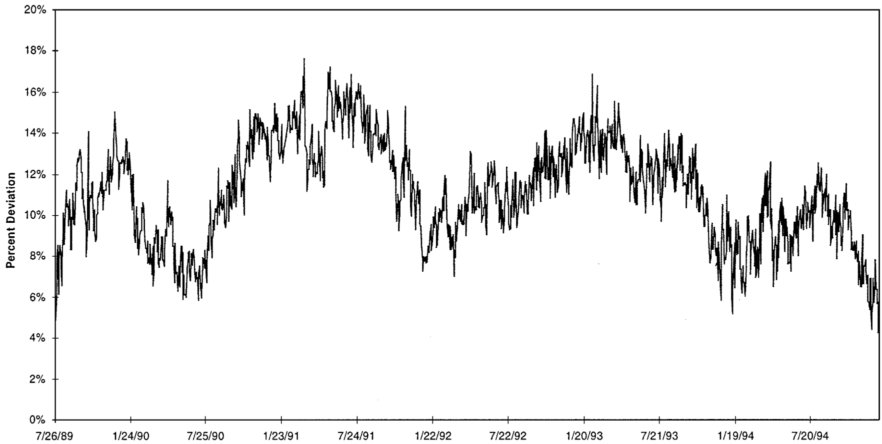

Fig. 3. Log deviations from SmithKline Beecham parity. Note: This figure shows on a percentage basis the deviations from theoretical parity of SmithKline Beecham H and E shares traded on the NYSE. Data are from the Center for Research in Security Pricing (CRSP).

where A and B represent the twin pair. Because of the cross-border aspects of these markets, we include currency changes as well as local-currency stock returns as market factors in Eq. (1). The null hypothesis is that all of the slope coefficients are zero. Under the alternative hypothesis, the more a stock trades on a given market, the higher its estimated slope. So for example, since Unilever N.V. trades relatively less intensively in the U.K. than Unilever PLC, the relative return of N.V. over PLC should generate a negative coefficient on the FTSE, and a positive coefficient on the S&P and Dutch markets (where N.V. trades relatively more intensively). Similarly, the N.V./PLC differential should exhibit a negative coefficient on the guilder/dollar and guilder/pound exchange rates. For given local-currency stock returns, an appreciation of the guilder increases the return on the Dutch index relative to other indexes, and therefore should increase the N.V./PLC differential.

Clearly, the log dollar return on a foreign stock index can be written as the sum of the local-currency stock return plus the log currency change. We use this additive decomposition to give each market and currency factor its own coefficient in Eq. (1), which is preferred to imposing the same coefficient for several reasons. First, currency values and local-currency stock prices are typically recorded at different times of day, inducing measurement error in the dollar returns. By separating out the two factors, we keep any measurement error in one of the variables from infecting the coefficient on the other currency change and local-market stock return are nearly uncorrelated). Second, any change in the dollar value of foreign stocks must be due to some combination of currency change and local stock return. It is useful to know if the relative twin returns have differential exposures to these two factors. For example, if local residents drive up the local-currency value of local stocks (caused by, say, a decline in risk aversion or by noise), they may drive up the price of the 'home' twin relative to the 'foreign' twin. We would therefore expect to find a positive beta on the appropriate local-currency stock index in Eq. (1). But, changes in the local currency may be driven by entirely different factors, so that the beta on the currency change could be zero.

The data in Table 1 suggest that under the alternative hypothesis, Royal Dutch should have higher correlation with the U.S. and Dutch markets, while Shell should have higher correlation with the U.K. market. The same is true for the relative returns on Unilever N.V. and PLC. For SmithKline Beecham, the A (or H) share/E share differential should vary positively with the U.S. market and negatively with the U.K. market.

We estimate Eq. (1) using return horizons of one, two, five, 15, and 50 days. The lower frequency regressions are less affected by imperfect synchronization of price observations (e.g., prices are observed at the closes of the New York and European markets, which occur with a five-hour difference), staleness, bid/ask bounce, etc. Furthermore, these tests can help differentiate among the underlying causes of segmentation. For example, if liquidity shocks explain the comovement of local market stocks, they should do so predominantly at higher frequencies.

We also examine the twin price differential for evidence of univariate mean reversion at very low frequencies. Specifically, we test to see whether we can reject the hypothesis that twin price disparities contain unit roots:

$$\Delta P_{\mathrm{A-B},t} = \alpha + \delta t + \beta P_{\mathrm{A-B},t-1} + \gamma(\Delta P_{\mathrm{A},t-1} - \Delta P_{\mathrm{B},t-1}) + \varepsilon_t, \tag{2}$$

where $P_{\mathrm{A-B},t}$ is the difference in the logs of twin prices, and Δ is the first-difference operator. The null hypothesis of a unit root in price differentials is given by $\beta = 0$. Naturally, this null hypothesis is unlikely to be true: it is hard to accept the notion that the price differential contains a unit root, so that over sufficient time, the probability that the differential becomes arbitrarily large equals one. However, we use Eq. (2) to get a point estimate of the rate at which price differentials decay. We also investigate the multivariate comovement of price disparities and market indexes. In particular, we test whether price disparities are cointegrated with some linear combination of stock indexes.

4. Data

European stock prices for Shell and Unilever PLC are taken from the London Stock Exchange, while the prices of Royal Dutch and Unilever N.V. are from the

Amsterdam Exchange.[7] Royal Dutch, Shell, Unilever PLC, and Unilever N.V. are traded as American Depository Receipts (ADRs) in the U.S. Royal Dutch trades in the U.S. market as a regular security.[8] U.S. return data are from Center for Research in Security Prices (CRSP). The sample period is January 1, 1980 to December 31, 1995. European prices for SmithKline Beecham A shares are from Interactive Data Corporation and dividend data are from Bloomberg Data Services. SmithKline Beecham E shares and ADRs of the A shares (H shares) are from CRSP. The sample period follows the merger of SmithKline and Beecham, July 26, 1989 to December 31, 1995. All returns are expressed in log form.

For U.S. and U.K. market returns, we use log returns of the S&P 500 and FTSE indexes, respectively. The use of these popular indexes creates some ambiguity because Royal Dutch and Unilever N.V. are in the S&P 500 and Shell, Unilever PLC, and SmithKline Beecham are in the FTSE. Consequently, the regression coefficients are slightly biased relative to what they would be on indexes which exclude these stocks. The bias is minor since these stocks comprise only a small part of index capitalization. To see this, one can estimate the approximate bias in the coefficient relative to what it would be in the absence of an own-stock effect. Using data on capitalizations, covariances and variances from 1994, e.g., we calculate an upward bias of 0.032 in the coefficient for Shell, which has the largest capitalization of the three stocks in the FTSE.[9] This source of bias is too small to affect the results presented below.[10]

The own-stock effect is more severe in the case of the Netherlands stock index. Royal Dutch is by far the largest native stock traded on the Amsterdam Exchange. To eliminate any confusion, we remove Royal Dutch from the standard CBS Allshare General Price index. Data for this index and all other European indexes and exchange rates are obtained from Datastream.

---

[7] Data for Royal Dutch, Shell and Unilever PLC are total returns from Datastream. For Unilever N.V., we use price data from Interactive Data Corporation, and total return data from Datastream (January 1, 1993 to December 31, 1995). We obtain dividend information for Unilever N.V. from Rosenthal and Young (January 1, 1980–May 16, 1986), corporate annual reports (May 17, 1986 to May 4, 1989), and Bloomberg (May 5, 1989 to December 31, 1992).

[8] Shell Oil U.S. handles shareholder servicing responsibilities for Royal Dutch in the U.S., making ADRs unnecessary.

[9] The bias in beta is given by

$$\beta_{w} - \beta_{w/o} = \left(\frac{Cov(r_{sh}, r_{ftse})}{Var(r_{ftse})}\right) - \left(\frac{Cov(r_{sh}, r) - \alpha Var(r_{sh})}{Var(r_{ftse}) - \alpha^2 Var(r_{sh}) - 2\alpha Cov(r_{sh}, r_{ftse})}\right),$$

where β_w and $\beta_{w/o}$ are regression coefficients with and without Shell included in the FTSE, and α is Shell's fraction of the FTSE's capitalization (equal to 0.030 in 1994). Using data from 1994 to estimate the variances and covariances above, β_w and $\beta_{w/o}$ are estimated as 0.913 and 0.891, respectively. This suggests that the beta estimate is approximately 0.02 too high.

[10] In some tests (not reported), we create our own value-weighted U.K. stock index of the 20 largest U.K. stocks (as of 1993) excluding Shell, Unilever PLC, and SmithKline Beecham. The coefficients on this index are nearly identical to those on the FTSE.

Another important consideration is where returns are measured. In the tables below, we estimate the relative return on the twins by taking the difference of their log returns *in the markets where they trade most actively*. For example, we use the returns of Royal Dutch and Shell in Amsterdam and London. The basic results are unaffected if we use instead the relative return of Royal Dutch and Shell observed in, say, New York. In other words, the results we report are not sensitive to geographic deviations in the law of one price for any given stock.

A final issue concerns the currency denomination of returns. We leave all return variables in local currencies and then add exchange-rate changes as separate independent variables on the right-hand side of the regressions. To the extent that exchange rates and local-currency equity returns are uncorrelated, any error in exchange-rate changes from non-synchroneities will not bias the coefficients.[11]

5. Results

5.1. Alternative specifications

Tables 2–4 report estimates of Eq. (1) for Royal Dutch/Shell, Unilever N.V./PLC, and SmithKline Beecham, respectively.[12] Each line in the tables represents a slight variant of the general specification of the regression. The first four specifications use one-day return horizons, while specifications 5–8 use longer return horizons. For the one-day returns, specifications 1 and 2 represent slightly different lead/lag variants. In specification 1, the independent variables have one lead and one lag of all right-hand side variables. In specification 2, we restrict the leads and lags to those suggested by the actual market timing differences. For example, in Table 2, the dependent variable, the relative return of Royal Dutch over Shell, is observed daily at the close of European trading. Since the European markets close before the U.S. market, only the earlier day's

[11] Exchange-rate changes and local currency stock returns show little correlation in our data. In an earlier version of this paper (available from the authors), we provide a second method of dealing with currencies. We convert all returns into a common currency, and omit exchange-rate changes from the right-hand side of the regressions. In principal this method is inferior, because non-synchronous measurement of currency rates and stock prices introduce measurement error into the right-hand side variables. However, in practice the two methodologies yield very similar results.

[12] In the tables, twin equity returns are observed in the country where each twin is most liquid. We tried using returns from a common market (e.g., Royal Dutch and Shell both measured on the NYSE). See the earlier version of this paper for details. The results were qualitatively similar to those presented here. Small differences in coefficients (particularly in the 1-day regressions) occur, however, due to transient deviations from the law of one price for any given stock.

Table 2
Royal Dutch/Shell price differentials and market movements

This table reports regression estimates of the equation:

$$r_{\mathrm{RD-SH},t} = \alpha + \sum_{i=-1}^{1} \beta_i S\&P_{t+i} + \sum_{j=-1}^{1} \delta_j FTSE_{t+j} + \sum_{k=-1}^{1} \lambda_k DI_{t+k} + \sum_{l=-1}^{1} \gamma_l gl/\$_{t+l} + \sum_{m=-1}^{1} \upsilon_m gl/£_{t+m} + \varepsilon_t$$

where $r_{\mathrm{RD-SH},t}$ is the difference between the log returns of Royal Dutch (Amsterdam) and Shell (London); *S&P*, *FTSE*, and *DI* are returns on the *S&P*, Financial Times Allshare index, and Dutch stock indexes, respectively, expressed in their native currencies; and *gl*/\$ and *gl*/£, represent log changes in the guilder-to-dollar and guilder-to-pound exchange rates. Specification 1 includes leads and lags (shown) to allow for nonsynchronous trading. Specification 2 employs a more restricted set of leads and lags (based on actual time differentials). Specifications 3 and 4 are the same as Specifications 1 and 2, but include a lagged dependent variable on the right-hand side. Durbin's Alternate H (DAH) is reported in place of the Durbin–Watson (DW) statistic for Specifications 3 and 4. Specifications 5, 6, 7, 8 employ 2-, 5-, 15-, and 50-day returns. For these specifications, leads and lags of independent variables are dropped. All regressions are OLS, with standard errors that allow for serial correlation and heteroskedasticity. Where there is only a single coefficient, standard errors are in parentheses.

Specification	Return horizon	R^2	DW or DAH	DOF	Lagged dep. var.	S&P	FTSE	Dutch index	*gl*/\$	*gl*/£
1, 1980–1995	1 day	0.247	2.37	4155		0.207[c]	– 0.428[c]	0.150[c]	– 0.102[c]	– 0.345[c]
2, 1980–1995	1 day	0.218	2.35	4164		0.135[c]	– 0.516[c]	0.365[c]	– 0.123[c]	– 0.612[c]
3, 1980–1995	1 day	0.271	– 0.39[c]	4154	– 0.174[c]	0.205[c]	– 0.516[c]	0.213[c]	– 0.113[c]	– 0.439[c]
4, 1980–1995	1 day	0.262	0.19	4164	– 0.209[c]	0.146[c]	– 0.536[c]	0.359[c]	– 0.121[c]	– 0.612[c]
5, 1980–1995	2 days	0.204	2.42	1950		0.064[b]	– 0.451[c]	0.292[c]	– 0.041[a]	– 0.502[c]
						(0.032)	(0.038)	(0.032)	(0.030)	(0.047)
6, 1980–1995	5 days	0.244	2.29	776		0.087[b]	– 0.409[c]	0.246[c]	– 0.068[a]	– 0.440[c]
						(0.038)	(0.042)	(0.041)	(0.046)	(0.070)
7, 1980–1995	15 days	0.233	2.49	254		0.116[c]	– 0.370[c]	0.213[c]	– 0.126[b]	– 0.287[c]
						(0.048)	(0.048)	(0.053)	(0.059)	(0.070)
8, 1980–1995	50 days	0.521	2.35	71		0.184[c]	– 0.489[c]	0.285[c]	– 0.170[c]	– 0.385[c]
						(0.078)	(0.066)	(0.060)	(0.072)	(0.102)

2, 1980	1 day	0.187	2.40	250	0.074	– 0.636[c]	0.450[c]	– 0.114	– 0.629[c]
2, 1981	1 day	0.274	2.33	253	0.483[b]	– 0.882[c]	0.817[c]	– 0.449[b]	0.885[c]
2, 1982	1 day	0.188	2.29	253	0.186	– 0.540[c]	0.356[c]	– 0.152	– 0.846[c]
2, 1983	1 day	0.265	2.05	253	0.291[c]	– 0.500[c]	0.141[a]	– 0.065	– 0.779[c]
2, 1984	1 day	0.305	2.19	253	0.206	– 0.556[c]	0.364[c]	0.024	– 0.752[c]
2, 1985	1 day	0.158	2.39	253	– 0.036	– 0.307[c]	0.158[b]	– 0.050	– 0.562[c]
2, 1986	1 day	0.295	2.02	253	0.131[c]	– 0.323[c]	0.198[c]	– 0.067	– 0.564[c]
2, 1987	1 day	0.293	2.38	253	0.048[c]	– 0.496[c]	0.484[c]	0.212	– 0.656[c]
2, 1988	1 day	0.270	2.69	253	0.084[b]	– 0.630[c]	0.437[c]	– 0.178	– 0.583[c]
2, 1989	1 day	0.362	2.16	253	0.069	– 0.722[c]	0.464[c]	– 0.177[b]	– 0.345[c]
2, 1990	1 day	0.256	2.43	253	0.091[b]	– 0.306[c]	0.247[c]	– 0.182	– 0.695[c]
2, 1991	1 day	0.189	2.09	253	0.033	– 0.562[c]	0.499[c]	– 0.005	– 0.328[a]
2, 1992	1 day	0.242	2.23	253	0.151	– 0.428[c]	0.289[c]	– 0.187[b]	– 0.430[c]
2, 1993	1 day	0.323	2.27	253	– 0.097	– 0.475[c]	0.266[c]	– 0.009	– 0.659[c]
2, 1994	1 day	0.376	2.45	253	0.224[c]	– 0.698[c]	0.388[c]	0.260[c]	– 0.556[c]
2, 1995	1 day	0.183	2.65	252	0.059	– 0.270[c]	0.186[b]	– 0.169[b]	– 0.357[c]

[a]Significant at the 10% level for F-tests that the sum of all coefficients (leads and lags) equals zero.
[b]Significant at the 5% level.
[c]Significant at the 1% level.

Table 3
Unilever N.V./Unilever PLC price differentials and market movements

This table reports regressions estimates of the equation:

$$r_{\mathrm{NV-PLC},t} = \alpha + \sum_{i=-1}^{1} \beta_i S\&P_{t+i} + \sum_{j=-1}^{1} \delta_j FTSE_{t+j} + \sum_{k=-1}^{1} \lambda_k DI_{t+k} + \sum_{l=-1}^{1} \gamma_l gl/\$_{t+l} + \sum_{m=-1}^{1} \upsilon_m gl/£_{t+m} + \varepsilon_t$$

where $r_{\mathrm{NV-PLC},t}$ is the difference between the log returns of Unilever N.V. (Amsterdam) and PLC (London); *S&P*, *FTSE*, and *DI* are returns on the S and P, Financial Times Allshare index, and Dutch stock indexes, respectively, expressed in their native currencies; and *gl*/\$ and *gl*/£, represent log changes in the guilder-to-dollar and guilder-to-pound exchange rates. Specification 1 includes leads and lags (shown above) to allow for nonsynchronous trading. Specification 2 employs a more restricted set of leads and lags (based on actual time differentials). Specifications 3 and 4 are the same as Specifications 1 and 2, but include a lagged dependent variable on the right-hand side. Durbin's Alternate H (DAH) is reported in place of the Durbin–Watson (DW) statistic for Specifications 3 and 4. Specifications 5, 6, 7, 8 employ 2-, 5-, 15-, and 50-day returns. For these specifications, leads and lags of independent variables are dropped. All regressions are OLS, with standard errors that allow for serial correlation and heteroskedasticity. Where there is only a single coefficient, standard errors are in parentheses.

Specification	Return horizon	R^2	DW or DAH	DOF	Lagged dep. var.	S&P	FTSE	Dutch index	*gl*/\$	*gl*/£
1, 1980–1995	1 day	0.290	2.27	4124		0.098[c]	– 0.490[c]	0.328[c]	– 0.138[c]	– 0.463[c]
2, 1980–1995	1 day	0.259	2.25	4133		0.046[c]	– 0.624[c]	0.556[c]	– 0.125[c]	– 0.658[c]
3, 1980–1995	1 day	0.298	– 0.30[b]	4091	– 0.131[c]	0.098[c]	– 0.571[c]	0.394[c]	– 0.157[c]	– 0.552[c]
4, 1980–1995	1 day	0.287	0.13[c]	4101	– 0.182[c]	0.085[c]	– 0.640[c]	0.544[c]	– 0.667[c]	– 0.132[c]
5, 1980–1995	2 days	0.258	2.26	1950		0.041[b]	– 0.550[c]	0.467[c]	– 0.090[c]	– 0.565[c]
						(0.024)	(0.033)	(0.032)	(0.033)	(0.052)
6, 1980–1995	5 days	0.244	2.26	776		0.034	– 0.470[c]	0.341[c]	– 0.123[c]	– 0.374[c]
						(0.042)	(0.044)	(0.048)	(0.039)	(0.072)
7, 1980–1995	15 days	0.239	2.44	254		0.095[b]	– 0.436[c]	0.253[c]	– 0.146[b]	– 0.291[c]
						(0.057)	(0.059)	(0.050)	(0.068)	(0.102)
8, 1980–1995	50 days	0.352	2.16	71		0.017	– 0.376[c]	0.274[c]	– 0.090[a]	– 0.255[c]
						(0.082)	(0.065)	(0.068)	(0.070)	(0.099)

2, 1980	1 day	0.300	2.06	247	0.073	– 0.596[c]	0.847[c]	– 0.401	– 0.862[c]
2, 1981	1 day	0.313	2.18	250	0.014	– 0.752[c]	0.760[c]	0.009	– 0.705[c]
2, 1982	1 day	0.331	2.03	250	– 0.092	– 0.687[c]	0.725[c]	– 0.145	– 0.777[c]
2, 1983	1 day	0.163	2.32	250	0.166	– 0.392[c]	0.247[c]	– 0.070	– 0.468[c]
2, 1984	1 day	0.355	2.29	250	– 0.007	– 0.546[c]	0.547[c]	– 0.120	– 0.755[c]
2, 1985	1 day	0.235	1.73	251	0.074	– 0.506[c]	0.390[c]	0.064	– 0.799[c]
2, 1986	1 day	0.355	2.05	251	– 0.010	– 0.442[c]	0.512[c]	– 0.417[c]	– 0.940[c]
2, 1987	1 day	0.291	2.34	251	– 0.060	– 0.744[c]	0.695[c]	– 0.093	– 0.886[c]
2, 1988	1 day	0.395	2.45	252	0.167[c]	– 0.778[c]	0.715[c]	0.101	– 0.510[c]
2, 1989	1 day	0.469	2.00	252	0.040[b]	– 0.696[c]	0.688[c]	– 0.214[a]	– 0.838[c]
2, 1990	1 day	0.346	2.21	250	0.188[a]	– 0.629[c]	0.548[c]	– 0.106	– 0.454[c]
2, 1991	1 day	0.256	2.16	250	0.080[a]	– 0.635[c]	0.502[c]	– 0.116	– 0.432[b]
2, 1992	1 day	0.220	2.21	251	0.199	– 0.369[c]	0.309[c]	– 0.127	– 0.450[c]
2, 1993	1 day	0.176	2.58	253	0.002[b]	– 0.493[c]	0.202[a]	– 0.069	– 0.688[c]
2, 1994	1 day	0.200	2.59	253	0.513[c]	– 0.775[c]	0.199	– 0.668[c]	– 0.845[c]
2, 1995	1 day	0.160	2.67	252	0.015[b]	– 0.456[c]	0.230[b]	– 0.213	– 0.464[b]

[a]Significant at the 10% level for F-tests that the sum of all coefficients (leads and lags) equals zero.
[b]Significant at the 5% level.
[c]Significant at the 1% level.

Table 4
SmithKline Beecham price differentials and market movements

This table reports regressions estimates of the equation:

$$r_{\mathrm{SKA-SKB},t} = \alpha + \sum_{i=-1}^{1} \beta_i S \wedge P_{t+i} + \sum_{j=-1}^{1} \delta_j FTSE_{t+j} + \sum_{l=-1}^{1} \gamma_l \$/£_{t+l} + \varepsilon_t$$

where $r_{\mathrm{SKA-SKB},t}$ is the difference between the log returns of SmithKline Beecham A shares (London) and E shares (New York); S&P and FTSE, are returns on the S&P and Financial Times Allshare index, respectively, expressed in their native currencies; and \$/£ represents log changes in the dollar-to-pound exchange rate. Specification 1 includes leads and lags (shown above) to allow for nonsynchronous trading. Specification 2 employs a more restricted set of leads and lags (based on actual time differentials). Specifications 3 and 4 are the same as Specifications 1 and 2, but include a lagged dependent variable on the right-hand side. Durbin's Alternate H (DAH) is reported in place of the Durbin–Watson (DW) statistic for Specifications 3 and 4. Specifications 5–8 employ 2-, 5-, 15-, and 50-day returns. For these specifications, leads and lags of independent variables are dropped. All regressions are OLS, with standard errors that allow for serial correlation and heteroskedasticity. Where there is only a single coefficient, standard errors are in parentheses.

Specification	Return horizon	R^2	DW or DAH	DOF	Lagged dep. var.	S&P	FTSE	\$/£
1, 7/89–12/95	1 day	0.221	2.70	1665		-0.270^{c}	0.291^{c}	0.119^{c}
2, 7/89–12/95	1 day	0.216	2.69	1668		-0.390^{c}	0.390^{c}	0.215^{c}
3, 7/89–12/95	1 day	0.311	-0.54^{c}	1665	-0.335^{c}	-0.508^{c}	0.458^{c}	0.212^{c}
4, 7/89–12/95	1 day	0.307	-0.43^{c}	1667	-0.318^{c}	-0.541^{c}	0.365^{c}	0.214^{c}
5, 7/89–12/95	2 days	0.118	2.70	834		-0.466^{c}	0.409^{c}	0.184^{c}
						(0.064)	(0.053)	(0.045)
6, 7/89–12/95	5 days	0.167	2.68	330		-0.460^{c}	0.380^{c}	0.136^{c}
						(0.069)	(0.055)	(0.051)
7, 7/89–12/95	15 days	0.112	2.57	106		-0.275^{c}	0.216^{c}	0.092^{a}
						(0.085)	(0.058)	(0.067)
8, 7/89–12/95	50 days	0.217	1.98	28		-0.299^{b}	0.120^{a}	-0.057
						(0.133)	(0.085)	(0.085)

2, 7/89–7/90	1 day	0.450	2.35	253	– 0.713[c]	0.629[c]	0.309[c]
2, 7/90–7/91	1 day	0.302	2.57	256	– 0.400[b]	0.242[c]	0.331[b]
2, 7/91–7/92	1 day	0.282	2.50	256	– 0.167[c]	0.213[c]	0.232[c]
2, 7/92–7/93	1 day	0.214	2.88	256	– 0.278[c]	0.544[c]	0.237
2, 7/93–7/94	1 day	0.122	2.85	256	– 0.235[c]	0.382[c]	– 0.137[a]
2, 7/94–7/95	1 day	0.113	2.57	256	– 0.060[c]	0.154[c]	0.104[c]
2, 7/95–2/95	1 day	0.143	2.41	107	– 0.457[c]	0.285[a]	0.035[c]

[a]Signifitblfn cant at the 10% level for F-tests that the sum of all tblfn coefficients (leads and lags) equals zero.
[b]Significant at the 5% level.
[c]Significant at the 1% level.

U.S. market return is included on the right-hand side of specification 2. Specifications 3 and 4 are analogous to specifications 1 and 2, except that a lagged dependent variable is added to the right-hand side. This allows us to estimate the short-run versus long-run effects of a change in the market indicators on the twin price disparity:[13]

$$r_{A-B,s} = \alpha + \theta r_{A-B,s-1} + \beta r_{S\&P,s} + \delta r_{FTSE,s} + \lambda r_{DI,s} + \gamma gl/\$_s + \upsilon gl/£_s + \varepsilon_{A-B,s}. \tag{3}$$

The coefficient β can be interpreted as the short-run response of the return differential to a shock to the S&P 500, and $\beta/(1-\theta)$ can be interpreted as the long-run response. If prices tend to revert toward parity, then we should find that long-run responses are smaller than short-run responses, so that $\theta < 0$.

Specifications 5–8 report results for return horizons of two, five, 15, and 50 days using specification 2. Because low power does not appear to be a problem at these horizons, we use non-overlapping returns to make inferences more reliable.[14]

5.2. Estimates

The results in Tables 2–4 strongly reject the perfect-integration hypothesis. The signs of virtually all coefficients line up with our alternative hypothesis, and most are significantly different from zero at the 1% level.[15] The estimates are also economically large. In Table 2, for example, the one-day Royal Dutch/Shell return differential yields coefficients of about 0.15 on the S&P, -0.50 on the FTSE, and 0.30 on the Dutch index. The coefficients on the exchange rate changes are also large, at -0.10 and -0.50 for the guilder/dollar and guilder/pound exchange rates. An one-percent appreciation of the guilder against the dollar and pound, respectively, increases the relative price of Royal Dutch over Shell by about 10 and 50 basis points. These coefficient values also imply that a 1% appreciation of the dollar relative to the pound increases the relative price of Royal Dutch over Shell by about 40 basis points.

[13] Leads and lags in (3) are identical to those in (1) for all variables other than the lagged dependent variable. They are omitted to keep the notation simple.

[14] Non-overlapping returns fail to utilize all the information in the data. However, they generate higher quality standard errors because the residuals are serially uncorrelated under the null hypothesis.

[15] The significance tests are F-tests on the sum of the lead, current, and lag coefficients for each index.

It is also interesting to note that much of the variation in return differentials (which have an average annualized standard deviation of about 17%) is explained by Eq. (1). The R^2s in Table 2 are surprisingly high, around 20% for one-day returns and up to 50% for longer-horizon returns.

The coefficient estimates appear reasonably stable over time. Interestingly, a large change in Shell ownership occurred in 1985, when U.S. holdings rose to 8% from under 1%. Table 2 suggests that this change in ownership was associated with a decline in the S&P coefficient, consistent with our alternative hypothesis. Specifications 3 and 4 yield estimates of the lagged-dependent variable coefficient, θ_{AB}, of about -0.2, which is strongly statistically significant. This implies that the short-horizon beta coefficients are about 20% greater than their long-horizon counterparts. While this estimate is not small economically, it suggests that the comovements we measure persist over longer return horizons.

Tables 3 and 4 reveal a similar story for Unilever N.V./PLC and SmithKline Beecham. We reject the null hypothesis in most cases at the 1% level.

These results provide evidence of comovement between relative twin prices and market indexes for both short and long horizons. The data actually reveal an even stronger finding: in our sample, we find no statistical evidence that the comovement is at all transient. Specifically, we cannot reject the hypotheses that: (1) the price differentials contain unit roots, and (2) the price differentials and stock indexes are cointegrated.

In Table 5 we investigate whether the price differentials contain unit roots using the augmented Dickey–Fuller test. The data cannot reject the unit root hypothesis for any of the twins. The estimates from the Dickey–Fuller test also give us a sense for the half-life of price deviations, as measured from daily data. With a coefficient on the lagged twin price differential of 0.004, the half-life of price deviations works out to be almost exactly one-half year. However, this estimate is imprecise, and we cannot reject the hypothesis that the half-life is infinite.

In addition, we test for cointegration between the twin price differentials and arbitrary linear combinations of market indexes. The data reject the null hypothesis of no cointegration for all three sets of twins.[16] This suggests that we would need a longer time series to make even the minimal claim that price differentials do not grow with stock markets differentials over the long run, but instead revert back toward zero.

The basic interpretation of these unit root tests is that price deviations and their relations with market variables are highly durable – so much so that we cannot detect evidence that the price deviations mean revert, or that the price

[16] To save space, we do not repeat the results here. See Froot and Dabora (1998) for details.

Table 5
Cointegration and unit root tests

Augmented Dickey-Fuller tests of log price differentials and log prices

Variable	Coefficient	P-value	Results
$P_{\mathrm{RD},t} - P_{\mathrm{Shell},t}$	– 0.0034	0.2926	Fail to reject unit root
$P_{\mathrm{UNV},t} - P_{\mathrm{Uplc},t}$	– 0.0042	0.8729	Fail to reject unit root
$P_{\mathrm{SKA},t} - P_{\mathrm{SKE},t}$	– 0.0052	0.6212	Fail to reject unit root
Dutch index	– 0.0002	0.9845	Fail to reject unit root
FTSE index	– 0.0006	0.4106	Fail to reject unit root
S&P index	– 0.0007	0.6735	Fail to reject unit root

Variables are relative log prices of twin stocks, e.g., $P_{\mathrm{RD},t} - P_{\mathrm{Shell},t}$ is the log price of Royal Dutch relative to that of Shell. Index variables are stock market total return indexes. Coefficients are estimates of β from the augmented Dickey–Fuller regression, $\Delta P_{\mathrm{A-B},t} = \alpha + \delta t + \beta P_{\mathrm{A-B},t-1} + \gamma(\Delta P_{\mathrm{A},t-1} - \Delta P_{\mathrm{B},t-1}) + \varepsilon_t$.

differentials do not follow differentials in market indexes. While we do not take the null hypotheses of these tests too literally, the tests do demonstrate the high degree of persistence in the twin price differentials.

6. Explaining the comovement of relative prices and market indexes

In this section we analyze several potential explanations for the price deviations and their comovements with market indexes. In order to conserve space, we focus on the largest twin pair, Royal Dutch/Shell, although similar results obtain for all three twin pairs. While each explanation could be a source of slippage between relative prices, it appears none can explain a meaningful fraction of the price differentials or comovement patterns.

6.1. Preliminary issue: the mechanics of splitting cashflow

The Royal Dutch/Shell Group splits net income in the proportion 60 : 40. The Group's charter includes an arrangement for offsetting corporate taxes across countries, so that the 60 : 40 split applies on an after-corporate-tax basis. This policy was tested in 1972 when the U.K. introduced a tax system aimed at eliminating double taxation of dividend income, the Advance Corporation Tax (ACT). ACT provided dividend holders an offset against corporate taxes on dividends. Specifically, under ACT shareholders received dividends plus a tax credit from the government. Over time, the tax credit has varied slightly, but has typically been about 20% of the *gross* dividend (dividend plus credit).

The Group's response to ACT was to split the value of the credit 60 : 40, thereby neutralizing the distributional effects of ACT.[17] To see how this works, note that any credit going to Royal Dutch shareholders must come through the company (since the U.K. government credits under ACT apply only to Shell shareholders). Thus, the Group pays more than 60% of distributed dividends to Royal Dutch shareholders. Inclusive of ACT, the precise split is 652 : 435 – still a 60 : 40 ratio – where the ACT credit is 8.7% (i.e., 20% of the Shell gross dividend of \$0.435). Sixty percent of the credit (\$0.052) goes to Royal Dutch shareholders, bringing their payment to \$0.652. The remaining \$0.348 (\$1.000 − \$0.652) goes to Shell shareholders. Thus, the Group's direct shareholder payments are split 652 : 348, but Shell shareholders also receive the 8.7% credit to bring their after-tax share to \$0.348 + \$0.087 = \$0.435.[18]

The larger point here is simply that Royal Dutch/Shell actively maintains its 60 : 40 policy, even intervening to offset asymmetries in the two countries' corporate-tax regimes.

6.2. Discretion in the use of dividend income

One possible explanation for the price behavior is that the parent companies do not pass dividends directly to shareholders, but instead invest a portion of the funds independently. If this is the case, we would expect parent company prices to deviate from the calculated expected price ratio as investment returns varied. However, this does not appear to be the case. The 1907 merger agreement specifies that the parent companies are not to make their own investments, and that they are to pass the dividends received directly along to shareholders.[19]

---

[17] The 1907 merger agreement anticipated that income taxes paid by parent companies on group dividends would have to be split 60 : 40. However, taxes on dividends paid by *shareholders* were not included. Because the ACT behaves both as a group tax on dividends and as a Shell shareholder credit, there was a dispute within the group companies as to whether Shell shareholders were entitled, in the spirit of the original merger agreement, to the entire ACT credit or only 40% of that credit. From the inception of the ACT in 1972, the group held to a 60 : 40 split of the ACT credit. In 1977, the group resolved the dispute by deciding that the 60 : 40 split would continue, but that Shell shareholders were to receive supplementary dividends of 15% of normal dividends for the 1977–1984 period, in consideration of their claims (January 13, 1977 press releases by parent companies).

[18] The split can be obtained as follows. Let a represent the fraction of distributed dividends received by Shell shareholders and b represent the after-tax-credit value per unit of distributed dividends. Royal Dutch shareholders must receive $0.6b = 1 - a$. Shell shareholders receive b augmented by their tax credit, $b = 1 + a\tau/(1-\tau)$, where τ is the corporate income tax rate. If $\tau = 0.20$, then $a = 0.348$.

[19] 'Royal Dutch Petroleum has no operations of its own and virtually the whole of its income derives from its 60% interest in the companies collectively known as the Royal Dutch/Shell Group of Companies … .' (Royal Dutch 1994 Annual Report). 'The Shell Transport and Trading Company, PLC has no operations of its own and virtually the whole of its income derives from its 40% interest in the companies collectively known as the Royal Dutch/Shell Group of Companies' (Shell Transport and Trading 1994 Annual Report).

However, neither company pays out *all* distributed group earnings as shareholder dividends. Both parents maintain a cash reserve account to promote ease in rounding and 'to provide a cushion against extreme currency fluctuations'.[20] (*Guidance Notes For Investors and Analysts:* 1994, p. 23.) The policy is to keep reserves low, but the size of the reserve cushion varies from year to year. Annual reports and company interviews suggest that the reserve account is invested either in cash at a bank or in the form of short-term deposits with a duration of less than three months. To see if the reserve is important, we can cumulate dividends in a common currency, adjusting for splits and short-term interest rates. This provides us with a crude measure of deviations from a common reserve investment policy. If reserve funds withheld by the parents are invested at riskless interest rates, then the ratio of cumulative dividends would be constant. In fact, the ratio of cumulated dividends did deviate from the 60 : 40 ratio, but only by a maximum of about 75 basis points (see Fig. 4). Such deviations are far too small to explain the magnitude and volatility of the price differentials. Nevertheless, Fig. 4 is interesting since cumulated dividends appear to be correlated with the price differential at low frequencies.

6.3. *Differences between the parent companies' expenditures*

Another potential explanation for the price disparities is that parent company expenses differ. If expenses deviated substantially from the 60 : 40 ratio, then the net receipts of shareholders would deviate as well. However, expense deviations from 60 : 40 are far too small to explain our findings. Differential expenses for 1993, for example, impact each share by approximately 6 basis points. A generous capitalization of these expense differentials would yield share price differentials of only about 1%.

6.4. *Voting rights*

Differences in corporate control might explain price disparities. Royal Dutch has a 60% share in both cash flows as well as voting power, so it could use this power to damage Shell shareholders interests.[21] Fluctuations in the value of

[20] 'As the amounts dealt with under the investment reserve have been, or will be, substantially reinvested by the companies concerned, it is not meaningful to provide for taxes on possible future distributions out of earnings retained by those companies; it is furthermore not practicable to estimate the full amount of the tax or the withholding tax element' (Royal Dutch Shell, 1994 Annual Report).

[21] The internal control of the companies is set up as follows. Each parent has its own independent management. The members of the Board of Managers of Royal Dutch and the Managing Directors of Shell are also Group Managing Directors. They maintain positions on the boards of the three Group Holding Companies. The ratio of members on this Group Board is 60 : 40.

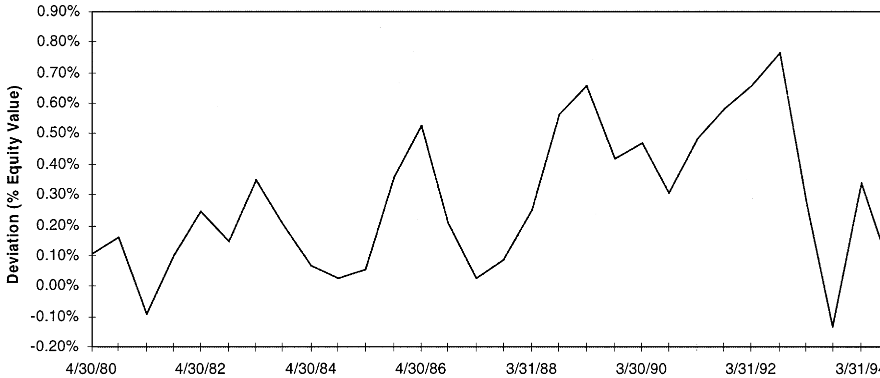

Fig. 4. Cumulative present value of dividends on Royal Dutch shares relative to those of Shell. Note: This figure shows, on a present value basis, the cumulative dividends for Royal Dutch relative to Shell as a percentage of the average stock price. Dividends of Royal Dutch are converted into a common currency and cumulated using short term interest rates.

control would lead to fluctuations in relative prices. The biggest problem with this story is that it fails to explain how Shell can be expensive relative to Royal Dutch, which was the case between 1980 and 1986. Furthermore, a control premium on Royal Dutch would explain the correlation with market indexes only if economy-wide changes in the value of control explain a large fraction of market movements. Finally, anti-takeover provisions make it difficult to accumulate large blocks of control of Royal Dutch or Shell. For example, ordinary shareholders of Royal Dutch face a cap on the number of shareholder votes at 12,000. This limits attacks on the management board, which can in principle alter the 60 : 40 relationship.

6.5. Dividends and currencies

Dividends are announced by both parents on the same day. At that time, dividend allocations for Royal Dutch (Shell) are converted into guilders (pounds) at prevailing spot exchange rates. In the time between the announcement and payment dates, fluctuations in the pound/guilder rate change the relative value of the dividend payments to Royal Dutch and Shell shareholders.

These factors can explain movements in the price differential, but only very minor ones. Exchange-rate changes matter only during the window between the announcement and ex-dividend dates. Furthermore, they can matter only for the value of the current dividend, not the present value of dividends. For example, assuming that the dividend/price ratio is 5%, dividend payments are made semi-annually, pound/guilder volatility is 1% per day, and actual payment periods corresponding to those in practice, currency differences in dividend denomination add at most 40 basis points to total return volatility over a year. This is very small relative to the large observed fluctuations in relative twin prices. Note also that we control for currency fluctuations in our regressions. Thus, currency fluctuations cannot explain comovements with local-currency market indexes.

6.6. Ex-dividend date structure

Royal Dutch and Shell shares can go ex-dividend on different dates. For example, during the 1991–1993 period, the difference between ex-dividend dates for Royal Dutch and Shell were 13 and 63 days, respectively, for interim and final dividend payments. This implies that, there will be a price wedge between the two securities if one security is past its ex-dividend date but the other is not. This effect is also small. At a dividend/price ratio of 5%, of which approximately 3% is the final dividend and 2% is the interim dividend, the price differential would be at most a few percent. There is also no reason to think that the ex-dividend patterns are correlated with market movements.

6.7. Tax-induced investor heterogeneity

Perhaps the most promising explanation for the price behavior is tax distortions. In the presence of such distortions, country-specific shocks to investor preferences or taxation could lead to correlation between relative twin returns and market indicators. However, for this explanation to succeed, taxes not only must segment one country from another, but within each country, taxes must also segment the twin pair.

To see this, suppose that there are differences in dividend taxation across countries and that, within any given country, dividends on twin stocks are treated identically by the local tax authority. Under these circumstances, a reduction in local dividend taxation might well move the local market up relative to the foreign market. However, there is no reason for the twin price differential to change, since from any given investor's perspective there is no change in the after-tax cashflows of one twin relative to the other. Thus, the tax treatment of one twin relative to the other must be different for at least some investor classes for the tax explanation to work.

To address this issue, we examine the tax burdens borne by specific investor groups in the U.S., U.K., and Netherlands. Taxation of international dividends is clearly complex. For example, a U.S. shareholder of a U.K. security might pay withholding tax, receive the ACT tax credit, and receive a credit from the U.S. Treasury on the withholding tax.[22] The actual rates paid may be altered

[22] This ignores taxes which affect both twins identically (e.g., personal income taxes).

through financial contracting or institutional restructuring. In spite of such complications, the tax laws are generally clear on how dividends ought to be treated for investor classes in different countries. Table 6 shows dividend withholding tax rates inclusive of ACT for shareholders by country and by investor class (private investors, companies and investment trusts, and pension funds).

The table shows that that private investors in all countries should be indifferent between investing in Royal Dutch and Shell.[23] Companies and investment trusts in the Netherlands and U.S. should also be indifferent between Royal Dutch and Shell, while U.K. companies and investment trusts should slightly prefer holding Shell. Pension funds, however, should not be indifferent between the twins. U.K. pension (or 'gross') funds pay no taxes on investments in Shell, but face 15% net withholding taxes on Royal Dutch dividends.[24] In contrast, Netherlands pension funds face no taxation on Royal Dutch, but pay 15% withholding taxes on Shell. Prior to January 1, 1994, U.S. pensions were indifferent to holding Royal Dutch and Shell, as they faced 15% withholding tax for both stocks. After January 1, 1994, the Double Taxation treaty between the U.S. and the Netherlands became effective, which gives U.S. pension funds a preference for Royal Dutch.

These facts have several implications. First, there is at least one group of investors in each country that is indifferent to the tax effect. This group could act as the marginal investor to equalize prices. For example, we expect private investors and companies in the Netherlands to hold shares in Shell when it is cheap relative to Royal Dutch. However, we find no discernible increase in the net holdings of Shell in the Netherlands during these periods.

Second, during all but the last two years of the sample period, *all* U.S. investors were indifferent to Royal Dutch and Shell on a tax basis. Thus, we expect to see holding patterns in the U.S. move toward the cheaper security. For example, Shell is relatively cheap from 1985 through 1992. Nevertheless, very few Shell shares are held in the U.S. during this period, yet at the same time Royal Dutch holdings in the U.S. are large and increasing. Furthermore, the tax

---

[23] When holding Royal Dutch, U.K. residents pay a 25% withholding tax, but 10% is reclaimable under the U.K./Netherlands double taxation agreement. The U.K. also levees a supplemental 5% dividend tax, bringing the total tax to 20%. The Shell shareholder also pays a net tax of 20% on dividends, so that the taxation on Royal Dutch and Shell are the same. Netherlands investors are subject to a 25% withholding tax on Royal Dutch dividends, which is creditable against their Netherlands income tax liability on the dividends. Shell shareholders that invest through a U.K. nominee company receive the full U.K. tax credit, but then must pay a 15% U.K. withholding tax. The withholding tax is creditable against Netherlands income taxes, so that the effective tax rates are equal on both sources of dividend income.

[24] Under U.K. law, tax-exempt investors, including pension funds in the U.S., U.K., and Netherlands, are entitled to a full credit against ACT.

Table 6
Taxation of different investor classes in different countries, 1993[a]

Country	Investor class	Tax rate on royal dutch dividends	Tax rate on shell dividends	Preference	Difference in annual return from tax differential[b]
UK	Private investors	20%	20%	Indifferent	–
	Companies	33%	20%	Shell	– 0.64%
	Pension funds	15%	–	Shell	– 0.74%
Nether	Private investors	25%	25%	Indifferent	–
	Companies	25%	25%	Indifferent	–
	Pension funds	–	25%	Royal Dutch	1.23%
US[c]	Private investors	15%	15%	Indifferent	–
	Companies	15%	15%	Indifferent	–
	Pension funds[d]	15%	15%	Indifferent	–

[a] Taxes represented withholding tax, dividend tax, and ACT. Tax treatment of capital gains on Royal Dutch and Shell were equivalent for all shareholder groups, and are therefore not reported.
[b] Average of Royal Dutch and Shell dividend/price ratios (4.92% in 1993) times the difference between Shell and Royal Dutch rates of dividend taxation.
[c] In the U.S., withholding taxes were reclaimable from income tax for corporations and individuals. Withholding taxes on foreign securities could either be deducted against U.S. personal or corporate income taxes, or, under current tax treaties, refunded directly from the U.K. and Netherlands tax authorities.
[d] Historically, U.S. pension and endowment funds were not able to deduct foreign taxes paid against U.S. tax obligations. Following January 1, 1994, U.S. pension funds were able to obtain withholding-tax refunds on Netherlands stocks, such as Royal Dutch, reducing the effective tax rate to zero.

indifference makes it difficult to explain the correlation of relative prices with either U.S. market returns or the U.S. dollar.

Third, even though some investors may have had tax-induced differences in reservation prices, it is not clear that these differences would be large enough to explain price deviations of 30% or more. Thus, tax issues, while potentially helpful, are unlikely to explain all of the components of the price deviations.

7. Conclusions

This paper presents evidence that stock prices are affected by the location of trade. It shows that twin stocks, which have nearly identical cashflows, move more like the markets where they trade most intensively than they should. The comovements between price differentials and market indexes are present at long as well as short horizons. Location of trade therefore appears to matter for pricing.

Our study suggests three possible sources of segmentation. The first source is tax-induced investor heterogeneity. This explanation seems incomplete. It does not explain correlations of twin price differentials with the U.S. market, since during the bulk of our sample all major U.S. investor groups faced equivalent tax treatment on twin stocks, and it does not explain why U.S. holdings of the cheap stock did not grow and why holdings of the expensive stock did not shrink.

The second possible source of segmentation is noise. Market-wide noise shocks from irrational traders, which infect locally traded stocks more than foreign traded stocks, can explain the comovements. Indeed, this story suggests that the portion of market movements that is correlated with fluctuations in twins' relative prices is attributable to noise. The main problem with this story – here and more generally – is that the source of noise or persistent irrationality is difficult to identify.

Third, institutional inefficiencies might explain comovements. By virtue of higher liquidity or inclusion in domestic-market indexes, one twin may be classified as a 'domestic' stock. (Note that causality here could easily run the other way, suggesting the possibility of multiple equilibria.) Classification as 'domestic' or 'foreign' appears to be important in practice, and could help resolve informational asymmetries and agency problems in the investment process.

Finally, there is the question of how arbitrage disciplines the price gap. In a frictionless world, it is clear that arbitrage would occur – any single investor could finance sufficiently large long positions to drive prices to parity.[25] But lack of disciplinary arbitrage does not explain why there are deviations in the first place.

References

Bodurtha, J., Kim, D.S., Lee, C., 1993. Closed-end country funds and U.S. market sentiment. University of Michigan and Sang Yong Research Institute.

Chen, N.F., Kan, R., Miller, M., 1993. Are the discounts on closed-end funds a sentiment index? Unpublished working paper. Journal of Finance 48, 795–800.

Chopra, N., Lee, C., Shleifer, A., Thaler, R., 1993. Yes, discounts on closed-end funds are a sentiment index. Journal of Finance 48, 801–808.

Froot, K.A., Dabora, E., 1998. How are stock prices affected by the location of trade. Working paper no. 6572. National Bureau of Economic Research, Cambridge, MA.

Froot, K.A., Perold, A., 1996. Global equity markets: the case of Royal Dutch and Shell. Case no. 296-077. Harvard Business School, Cambridge, MA.

Hardouvelis, G., La Porta, R., Wizman, T., 1995. What moves the discount on country equity funds. Working paper no. 4571. National Bureau of Economic Research, Cambridge, MA.

[25] Specific data on transactions costs and strategies are explored by Froot and Perold (1996).

Lee, C., Shleifer, A., Thaler, R., 1991. Investor sentiment and the closed-end fund puzzle. Journal of Finance 46, 75–109.

Pontiff, J., 1993. Costly arbitrage and closed-end fund discounts. Unpublished working paper. University of Washington, Seattle, WA.

Rosenthal, L., Young, C., 1990. The seemingly anomalous price behavior of royal dutch shell and unilever N. V./PLC. Journal of Financial Economics 26, 123–141.

Royal Dutch Petroleum Company, Shell Transport and Trading, 1994. Guidance Notes for Investors and Analysts.