

Akram, Q. Farooq; Rime, Dagfinn; Sarno, Lucio

Working Paper

Arbitrage in the Foreign Exchange Market: Turning on the Microscope

Working Paper, No. 2005/12

Provided in Cooperation with:

Norges Bank, Oslo

Suggested Citation: Akram, Q. Farooq; Rime, Dagfinn; Sarno, Lucio (2007) : Arbitrage in the Foreign Exchange Market: Turning on the Microscope, Working Paper, No. 2005/12, ISBN 82-7553-328-7, Norges Bank, Oslo,
<https://hdl.handle.net/11250/2498432>

This Version is available at:

<https://hdl.handle.net/10419/209859>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc-nd/4.0/deed.no>

Arbitrage in the Foreign Exchange Market: Turning on the Microscope*

Q. Farooq Akram^{a†}, Dagfinn Rime^a and Lucio Sarno^b

a: Norges Bank

b: University of Warwick, CEPR and AXA Investment Managers

This revision: February 2007

Abstract

This paper provides real-time evidence on the frequency, size and duration of arbitrage opportunities and deviations from the law of one price (LOP) in the foreign exchange market. We investigate deviations from the covered interest rate parity (CIP) condition ('round-trip arbitrage') and inter-market price differentials ('one-way arbitrage') using a unique data set for three major capital and foreign exchange markets that covers a period of more than seven months at tick frequency. The analysis unveils that: i) numerous short-lived violations of CIP and the LOP arise; ii) the size of CIP violations can be economically significant across exchange rates; iii) their duration is, on average, high enough to allow agents to exploit them, but low enough to explain why such opportunities have gone undetected in much previous research using data at lower frequency.

JEL classification: F31; F41; G14; G15.

Keywords: exchange rates; arbitrage; covered interest rate parity; foreign exchange microstructure.

***Acknowledgments:** This work was partly undertaken when Lucio Sarno was a Visiting Scholar at the International Monetary Fund and Norges Bank. The authors are grateful to a number of people who kindly offered constructive comments at various stages of this paper, including Ulf Axelson, Magnus Dahlquist, Martin Evans, Rich Lyons, Michael Melvin, Kristian Miltersen, Michael Moore, Carol Osler, Charlotte Østergaard, Mark Salmon, Per Strömberg, Mark Taylor, Paolo Vitale, in addition to participants in presentations at the 2006 European Finance Association Annual Meeting, Zurich; the 2006 European Economic Association Annual Meeting, Vienna; the 2005 Econometric Society World Congress, London; the 2005 Foreign Exchange Microstructure Conference, Aix-en-Provence; the 2005 High-Frequency Econometrics and Microstructure Workshop, University of Warwick; the Stockholm Institute for Financial Research; the University of Oslo; and Norges Bank. The authors also thank several foreign exchange traders and colleagues in the departments for market operations and interbank settlements at Norges Bank for insightful conversations, especially Jan Engebretsen and Thor Kristiansen. Lucio Sarno gratefully acknowledges financial support from the Institute for Quantitative Investment and Research (INQUIRE). The authors alone are responsible for the views expressed in the paper and for any errors that may remain.

[†]Corresponding author: Norges Bank, P.O. Box 1179 Sentrum, 0107 Oslo, Norway. Email: farooq.akram@norges-bank.no

1 Introduction

Arbitrage is one of the fundamental pillars of financial economics. It seems generally accepted that financial markets do not offer risk-free arbitrage opportunities, at least when allowance is made for transaction costs. This notion is also directly related to the law of one price (LOP), which postulates that in well-functioning efficient financial markets, identical securities must have the same price, no matter how they are created. For example, if a derivative instrument can be created using two different sets of underlying securities, then the total price for each derivative instrument must be the same or else an arbitrage opportunity would exist. Arbitrage is the mechanism that should ensure the validity of the LOP.

On the one hand, if all the market participants in financial markets post arbitrage-free tradable quotes, then arbitrage opportunities cannot exist. On the other hand, while the assumption of no arbitrage is likely to be reasonably mild or valid in several contexts in finance, its violations can be rationalized on several grounds. In general terms, the absence of arbitrage opportunities gives rise to the so-called ‘arbitrage paradox’, first pointed out by Grossman and Stiglitz (1976, 1980). That is, if arbitrage is never observed, market participants may not have sufficient incentives to watch the market, in which case arbitrage opportunities could arise. A possible resolution to this paradox is for very short-term arbitrage opportunities to arise, inviting traders to exploit them, and hence be quickly eliminated. Also, microstructure theory shows how price differences may occur for identical assets in markets that are less than fully centralized, segmented or with an imperfect degree of transparency (O’Hara, 1995; Lyons, 2001).¹

Empirical studies have, however, been unable to detect short-term arbitrage opportunities in a variety of financial markets. Given the high activity level in major financial markets, such short-term arbitrage opportunities can only be adequately studied using real-time quotations on all asset prices involved. Such data are, however, notoriously difficult to obtain. Furthermore, one must take into account all relevant aspects of the microstructure of the markets in order to capture the opportunities and transaction costs that market participants face.

This paper investigates empirically the existence of arbitrage and the properties of potential departures from no-arbitrage conditions using a microstructure perspective. Specifically, we study the foreign exchange (FX) market, for which no-arbitrage conditions are well known and relatively easy to test. These conditions are covered interest rate parity (CIP) and the related concept of the LOP for lending and borrowing services which has been termed ‘one-way arbitrage’ in the relevant literature.

¹See also the theories related to limits to arbitrage (Shleifer and Vishny, 1997; Lyons, 2001).

The CIP condition states that net returns on an investment that borrows at home and lends abroad (or vice versa) in similar interest-bearing assets will be zero when exchange rate risk is hedged through forward or swap contracts. The CIP condition, which can be denoted as the condition for absence of ‘round-trip arbitrage’ opportunities, is the cornerstone riskless no-arbitrage condition in the FX market. The LOP for lending and borrowing services requires that the domestic lending (borrowing) interest rate should be the same as the foreign lending (borrowing) interest rate when the latter is adjusted to fully hedge for exchange rate risk. Following Deardoff (1979), a number of studies refer to this LOP condition to describe the absence of ‘one-way arbitrage’, in the form of ‘owner arbitrage’ and ‘borrower arbitrage’.

Often, the distinction between the CIP condition and the LOP condition is not duly appreciated. The LOP condition, or absence of ‘one-way arbitrage’, is fundamentally different from the CIP condition, and, more generally, from the notion of pure arbitrage, i.e. instantaneous riskless profit net of all costs. This is because violations of the LOP do not imply riskless profits, but imply that the same need–desire to lend (owner arbitrage) or desire to borrow (borrower arbitrage)—can be met at two different prices at a point in time. It may therefore be more appropriate to consider violations of the LOP as a form of ‘mispricing’ rather than pure arbitrage opportunities, which do not require any initial endowment or borrowing need. The LOP for lending and borrowing services has not been tested as widely as the CIP condition, especially as a distinct concept from the CIP condition. However, while not informative about the existence of pure arbitrage opportunities, analysis of the LOP (one-way arbitrage) is nevertheless informative about the ability of financial markets to price efficiently the same asset.²

The key advantages of this study relative to all previous empirical analyses of arbitrage are our data set, and a precise account of transaction costs and pricing and trading conventions. A rigorous empirical examination of no-arbitrage interest rate conditions in the FX market places stringent requirements on the data used. Contemporaneous, tradable (firm) quotes of comparable domestic and foreign interest rates and spot and forward exchange rates are needed in order to establish whether an apparent deviation from no-arbitrage conditions actually represented a profitable opportunity to agents at a given time or not. Moreover, the high level of activity in FX and international capital

²Since Deardoff (1979) introduced the concept of one-way arbitrage, researchers often interpreted this concept as closely related to pure arbitrage, and as a stricter condition which would yield more violations of parity conditions, i.e. a more stringent test of no-arbitrage conditions than CIP. Taylor (1987) offered tests of one-way arbitrage as additional evidence that CIP is not violated. In this paper, we shall use the term ‘one-way arbitrage’ for consistency with the literature since Deardoff (1979). However, it is important to emphasize that violations of one-way arbitrage have a different economic interpretation from violations of CIP in that the former do not imply riskless profits and hence do not correspond to pure arbitrage. For a further discussion of these issues, see Section 2.B.

markets demands use of high-frequency, real-time quotes to characterize the properties of arbitrage opportunities, especially their duration. Finally, it is equally important to have a sufficiently long sample to draw general conclusions. Our data set is the first data set in this literature that possesses these characteristics to a large extent.

Empirical studies of arbitrage in the FX market so far have not employed data sets that meet the above-noted requirements, mainly because such data sets have been unavailable to researchers. These studies suggest that arbitrage opportunities do not generally arise in the FX market and mispricing is negligible when one accounts for estimated transaction costs.³ The move to electronic trading platforms in the 1990s, however, has made it possible to obtain long data samples of real-time quotations for rigorous empirical work. The move itself provides a motivation for a fresh analysis of arbitrage opportunities because of changes in trading practices and market characteristics induced by electronic platforms. The growing literature on high-frequency exchange rate behavior and FX market microstructure has not—to the best of our knowledge—studied arbitrage, focusing instead on a variety of other issues relating to international currency patterns, trading behavior, and the role of order flow in explaining exchange rate movements (e.g. Lyons, 1995, 2001; Osler, 2000, 2003, 2005; Covrig and Melvin, 2002; Evans, 2002; Evans and Lyons, 2002, 2005; Payne, 2003; Bjørnnes and Rime, 2005; Lyons and Moore, 2005).

Use of real-time quotations can also shed light on the validity of another proposed resolution of the arbitrage paradox, which is the anecdote that providers of interest rate and exchange rate quotes set their quotes such that they knowingly do not misprice or offer counterparts riskless profit opportunities—i.e. set prices that violate the LOP or the CIP condition. For example, if quotes are always set such that no-arbitrage conditions are ensured conditional on the latest quotes of other instruments, these conditions will hold continuously without requiring trade to actually take place. However, from the microstructure literature we know that prices can temporarily deviate from no-arbitrage values due to, for example, portfolio management. The presence of profitable arbitrage opportunities in real-time could also point towards possible inefficiencies in information gathering, profit maximization and/or other constraints on the part of quote providers.

Our data set includes contemporaneous tick quotes of exchange rates and interest rates that pertain to the most liquid segments of the FX and capital markets. The sample includes ask and bid quotes for three major US dollar spot exchange rates: euro, UK sterling and Japanese yen. It also includes ask and bid quotes for exchange rate swaps

³Studies of FX arbitrage—primarily based on tests of CIP—include Branson (1969), Frenkel (1973), Frenkel and Levich, (1975, 1977), Callier (1981), Taylor (1987, 1989), Clinton (1988), Rhee and Chang (1992), Fletcher and Taylor (1993), Aliber, Chowdhry and Yan (2003), and Juhl, Miles and Weidenmier (2006). We briefly review this literature in the next section.

and for interest rates on deposits in quoting and base currencies. The tick quotes cover a period of more than seven months spanning from February 13 to September 30, 2004, and is the longest and highest-frequency data set ever used for examining FX arbitrage. The data have been collected through Reuters trading system on special order.

To anticipate our key results, we find that trading aimed at exploiting no-arbitrage conditions is, on average, not profit-making. We provide evidence that there are numerous short-lived profitable deviations from both the CIP and the LOP for lending and borrowing services. The size of the profitable deviations is economically significant across exchange rates and comparable across different maturities of the interest rates examined. Their duration is, on average, high enough to allow agents to exploit these opportunities, but low enough to explain why such CIP and LOP violations have gone undetected in much previous research using data at lower frequencies. We find little evidence in favor of the view that prices for spot and forward rates and for money market instruments are set directly from the formulas of no-arbitrage conditions in real time. Finally, our results suggest that frequency, size and duration of apparent arbitrage opportunities decline with the pace of markets. Overall, our evidence is consistent with the Grossman-Stiglitz view of financial markets, where efficiency is not interpreted as a statement about prices being correct at each point in time but the notion that in efficiently-functioning financial markets very short-term arbitrage opportunities can arise and invite traders to exploit them, which makes it worthwhile to watch the relevant markets. This is the arbitrage mechanism that restores the arbitrage-free prices we observe on average.

The paper is organized as follows. Section 2 presents the concepts of the CIP and the LOP, i.e. of round-trip and one-way arbitrage, respectively, in the FX market. It also presents a brief review of the literature. For the sake of consistency with existing literature on arbitrage in the FX markets, we use the terms ‘one-way arbitrage’, ‘owner arbitrage’ and ‘borrower arbitrage’ interchangeably with (perhaps the more appropriate terms) ‘LOP’, ‘LOP for lending services’ and ‘LOP for borrowing services’, respectively. Section 3 briefly discusses quoting conventions, transaction costs and their implications for calculations of gains and losses from arbitrage of the different forms. In addition, this section describes the data set. Section 4 presents the main empirical findings, relating to frequency, size and durations of returns from arbitrage opportunities. Section 5 reports the results from the sensitivity analysis of the core results, and an analysis of whether and how characteristics of profitable arbitrage opportunities vary with market pace and market volatility. Section 6 briefly summarizes the main conclusions. Finally, the appendix presents further details on quoting conventions, calculations of days to maturity and transaction costs for different exchange rates and traded volumes.

tunity in the reverse direction. That is, the validity of condition (2) does not necessarily imply violation of condition (3), or vice versa.

2.B One-way Arbitrage: The LOP

Recognition of the fact that funds are borrowed and lent at different rates makes it important to consider the behavior of traders who are looking for the highest return on their endowments and of those who are looking for the cheapest borrowing opportunities. The concept of ‘owner arbitrage’ (OA) refers to the case where a trader has an endowment of funds in some currency and wants to lend the funds to obtain the highest possible net return. Such traders weigh the option of lending own funds at the market bid rate for the endowment currency, against the option of converting the funds to another currency at the spot exchange rate and lending them at the market bid rate for that currency, while eliminating the exchange rate risk at the maturity of the lending contract through a forward contract. In contrast, the concept of ‘borrower arbitrage’ (BA) refers to the case where a trader aims to finance an investment in the cheapest way and thereby gain by minimizing funding costs. Such traders face the option of borrowing funds in the desired currency directly, or to borrow funds in another currency and convert them to the desired currency at the spot exchange rate, while eliminating the exchange rate risk at the maturity of the borrowing contract through a forward contract.

It follows that any gain by lending or borrowing in one currency relative to another would imply a violation of the LOP, which states that identical securities must have the same price, irrespective of how they are created. However, as shown below, such a violation would not imply profitable round-trip arbitrage, i.e. profitable deviations from the CIP, requiring no initial funds or borrowing needs. In the following sub-sections we consider the two relevant cases of one-way arbitrage in the FX market and point out their relationships with the conditions for profitable round-trip arbitrage, i.e. the case of CIP.

2.B.1 Owner Arbitrage Definition

The LOP will prevent OA opportunities under the following conditions:

$$(1 + i_d^b) \geq 1/S^a(1 + i_f^b)F^b, \quad (4)$$

$$(1 + i_f^b) \geq S^b(1 + i_d^b)/F^a. \quad (5)$$

The left-hand-side elements in these inequalities are lower than those in the case of CIP given in conditions (2)–(3), and hence they may be violated more easily, and more often, than the CIP inequalities; see Panels I-II in Table 1. Consequently, absence of profitable OA opportunities is a stronger test than absence of profitable covered arbitrage

opportunities.⁶

It follows that profitable arbitrage in the case of CIP implies the existence of OA opportunities while the converse may not be true—i.e. violation of conditions (2) and (3) is a sufficient but not a necessary condition for the violation of conditions (4) and (5) respectively. Moreover, if OA opportunities are not present, neither will CIP arbitrage be profitable while the opposite may not be the case—i.e. validity of conditions (4) and (5) is a sufficient but not a necessary condition for the validity of conditions (2) and (3) respectively. In Table 1 (Panels I-II), we summarize the relationship between CIP and OA.

2.B.2 Borrower Arbitrage Definition

The LOP will prevent BA opportunities under the following conditions:

$$(1 + i_d^a) \leq 1/S^b(1 + i_f^a)F^a, \quad (6)$$

$$(1 + i_f^a) \leq S^a(1 + i_d^a)/F^b. \quad (7)$$

The first inequality is relevant when the funds are required in domestic currency, while the second one becomes relevant when funds are required in the foreign currency.⁷

Note that when borrowing another currency than the currency eventually desired, the trader must consider how much she must borrow, and then sell, of that currency to obtain one unit of the desired currency. For instance, since the borrowed amount of foreign currency must be converted to the domestic currency at the spot bid rate, a trader must borrow $1/S^b$ of the foreign currency to obtain one unit of domestic currency. At maturity, her debt in foreign currency will be $1/S^b(1 + i_f^a)$, but $1/S^b(1 + i_f^a)F^a$ in domestic currency would be required if she enters a forward contract of that maturity to buy the foreign currency in order to settle her debt.

⁶Alternatively, one may be interested in measuring the risk-free net return in another currency than the endowment currency. In that case, if the endowment is in the domestic currency, one would weigh the option of lending it in the domestic market and converting the resulting amount at maturity to foreign currency at the forward exchange rate, against the option of converting the endowment right away to the foreign currency at the spot exchange rate and lending the resulting amount in the foreign capital market. In this case, if the domestic currency is the quoting currency, one would face the forward exchange rate at the ask in the first option, and the spot exchange rate at the ask in the second option. However, if the domestic currency was the base currency, one would be facing the bid side of both the forward and the spot exchange rates. We do not consider these alternatives in the empirical work below.

⁷Here we implicitly assume that revenues used to serve the borrowing costs flow in the same currency as that for the funds required. However, revenues used to serve the borrowing cost may flow in a different currency than that of the funds required. Take the following example: the agent needs funds to cover some costs in domestic currency, but the revenues used to serve the borrowing costs in domestic currency flow in the foreign currency. Then, the agent would weigh the option of borrowing funds in the domestic market and converting the borrowing costs at maturity at the forward exchange rate, against that of borrowing in the foreign market and converting the borrowed amount at the spot exchange rate. If the domestic currency is the quoting currency, then one would have to sell the foreign currency (at the bid rate) forward as well as spot. In contrast, if the domestic currency is the base currency, one would need to buy the domestic currency (at the ask rate) spot as well as forward.

It appears that profitable CIP arbitrage when measured, e.g., from the viewpoint of a domestic arbitrageur precludes profitable BA opportunities for a domestic fund raiser while the converse may not be true. That is, violation of conditions (2) and (3) is a sufficient but not a necessary condition for the validity of conditions (6) and (7) respectively; see Panel III of Table 1. In other words, if a positive return can be gained in domestic currency by borrowing domestic funds to lend abroad, it will also be relatively dearer to borrow funds abroad (when measured in domestic currency), but the converse may not be the case. However, profitable CIP arbitrage when measured in domestic currency d implies that it will be profitable for a foreign investor to borrow domestic currency funds, convert them to the foreign currency at the spot rate, while covering the exchange rate risk at maturity through a forward contract. In other words, profitable CIP arbitrage from the viewpoint of a domestic (foreign) dealer implies BA from the viewpoint of a foreign (domestic) dealer; see Panel IV of Table 1.

2.B.3 Refining the Interpretation of CIP, OA and BA

The above results suggests that one-way arbitrage opportunities, i.e. deviations from the LOP, may occur more often than round-trip (pure) arbitrage opportunities. Another reason for this to be the case is due to the fact that one-way arbitrage requires excess supply of (desire to lend) or demand of (desire to borrow) funds, while round-trip arbitrage requires no own funds or borrowing needs. In general, an infinite price elasticity of supply of and/or demand for funds can instantaneously eliminate deviations from the LOP. In the present context, however, differences between the price of a money market (MM) security and that of a derived/synthetic security, depending on two separate security markets, FX and MM, may arise frequently since the two markets will be governed by own supply and demand conditions (liquidity), own speed (relatively faster in the FX market), and own trading and pricing practices. These considerations are important when interpreting and comparing potential deviations from CIP with OA and BA opportunities in empirical work.

The key difference between CIP violations and OA/BA violations is that the latter do not imply the existence of pure arbitrage opportunities, only the existence of non-zero inter-market price differentials that cannot be exploited to generate riskless profits. In this sense, one-way arbitrage is a relative value statement. This is in stark contrast with CIP violations, which ensure riskless profit using no capital. It seems possible, and indeed plausible, that in two segmented markets for the same asset a bid (ask) in one market can be higher (lower) than the bid (ask) in the other, as long as the best bid is not higher than the best ask.

2.C A Brief Review of the Literature

A landmark study in the literature on testing no-arbitrage conditions in the FX market is Taylor (1987), which questioned the published evidence of deviations from CIP as it was not based on contemporaneously sampled real-time quotes of comparable domestic and foreign interest rates and spot and forward exchange rates. Accordingly, it was not possible to know whether an apparent deviation from CIP actually represented a profitable opportunity to agents at a given time or not. Using data that closely met the strict requirements and formulas for calculating profits from CIP arbitrage used by dealers, Taylor (1987) found rather strong evidence of CIP, not observing a single profitable round-trip arbitrage opportunity, though he noted several profitable OA opportunities. This study employed interest rate and exchange rate data points that were recorded within approximately one minute of each other, by phoning several London-brokers at ten minutes frequency during the most active hours (9.00–16.30) over three days in 1985. Subsequent studies have not convincingly overturned the support for CIP provided by Taylor (1987); either because the data employed have been of lower frequency and/or profitable deviations from CIP have been estimated or calculated imprecisely. These studies include, *inter alia*, Rhee and Chang (1992) and Fletcher and Taylor (1993).

Still, the robustness of Taylor’s (1987) findings may be questioned on several grounds. First, the data spans a period that may be too short for inferring general conclusions. Second, the recorded quotes were not strictly contemporaneous since quotes could change during a minute. And third, the ten minute frequency at which the observations were recorded seems to be relatively low and do not enable one to adequately characterize dynamics of possible deviations from CIP arbitrage, which may contribute to resolve the Grossman-Stiglitz arbitrage paradox. The ten-minute interval frequency used by Taylor (1987) could, however, have been sufficiently high to provide accurate results using data from the mid-1980s than using more recent data, especially given that there was no centralized (electronic) market at that time.

The literature on testing no-arbitrage conditions in currency markets has been somewhat dormant in the last ten years or so. This is primarily because in the non-electronic, highly decentralized markets studied until the early 1990s, it would have been extremely difficult to improve on the quality of the data beyond Taylor’s (1987, 1989) papers. Subsequent studies generally supporting the absence of arbitrage opportunities includes Rhee and Chang (1992), Fletcher and Taylor (1993), Aliber, Chowdhry and Yan (2003), and Juhl, Miles and Weidenmier (2006).

3 Data and Calculations of Returns from Arbitrage

We obtained data, on special order, from the Reuters trading system, which embeds general market quoting and maturity conventions. In this section, we present precise formulas for calculating deviations from the different no-arbitrage conditions in light of these conventions as well as transaction costs that a trader would typically face when dealing through this system. Appendices A.A and A.B provide a detailed account of quoting conventions, calculations of days to maturity and transaction costs for different exchange rates and traded volumes.

3.A Formulas Used for the Calculations

In the interbank-market dealers trade swaps rather than (outright) forwards. Swaps are denominated in so-called swap points, which express a multiple of the difference between forward and spot exchange rates. By convention, all of the spot exchange rates are quoted with four decimals, except for the Japanese yen, where two decimals are used. The smallest measure of movement for an exchange rate is called a “pip”. Swap points, which are expressed in pips, are therefore obtained by multiplying the difference between forward and spot exchange rates by 10^4 in general, and by 10^2 in the case of the Japanese yen.

We investigate potential returns from arbitrage by comparing the swap points quoted through Reuters with corresponding derived (or theoretical) swap points. The derived points can be obtained by rewriting the formulas presented above, (2)–(7), while taking into account relevant quoting and maturity conventions. For example, the deviation from the CIP on the bid side, condition (2), can be expressed as:

$$Dev_{CIP}^b = (F^b - S^a) - \frac{S^a(i_d^a \times \frac{D}{360} - i_f^b \times \frac{D}{360})}{(100 + i_f^b \times \frac{D}{360})} \times 10^4, \quad (8)$$

where the first right-hand term denotes market swap points for a given maturity obtained from Reuters, while the second term represents the corresponding derived swap points. In order to calculate derived swap points that are directly comparable to market swap points quoted on Reuters, we adjust the interest rates, which are quoted in percent per annum, to obtain interest rates for maturities less than a year.

Specifically, D denotes the number of days to maturity of swap and deposit contracts. It is calculated as the actual number of business days between the (spot) value date and the maturity date of a contract while taking into account bank holidays in the home countries of currencies and securities, and other conventions—see Appendix A.A for details. In general, the total number of days to maturity in a year are 360. For sterling contracts, however, the total number of days in a year are set at 365 in line with market conventions.

Thereafter, the resulting term is multiplied by 10^4 (or 10^2 in the case of the Japanese yen) to obtain the derived swap points. Deviations from a no-arbitrage condition, e.g. (8), are expressed in pips since they are defined as the difference between quoted and derived swap points.

In our empirical analysis, we treat the quoting currency as the domestic currency (d) and the base currency as the foreign currency (f), for convenience, since we overlook cases where both the quoting as well the base currencies are actually foreign currencies for a dealer. Table 2 makes explicit the quoting and base currencies for the three exchange rates examined.

Table 3 presents derived deviations for all of the no-arbitrage conditions pertaining to CIP, OA and BA. Deviations are profitable if equation (8), or the equivalent equation from Table 3, is positive net of other transactions costs. That is, in actual calculations of returns from arbitrage, we deduct 1/10 of a pip (10^{-5}) from the expressions for returns presented in Table 3 to obtain returns less brokerage and settlement costs. Appendix A.B shows that the sum of brokerage and settlement costs are at most 1/10 of a pip of a US-dollar pip for an arbitrage deal of required size.⁸ Thus, the number and size of profitable returns obtained by us are likely to represent lower bounds on the number of profitable returns through arbitrage.

3.B Data

We employ tick data collected via a continuous feed from Reuters over the period February 13–September 30, 2004. The data set allows us to investigate both round-trip as well as one-way arbitrage for three major exchange rates at four different maturities: 1, 3, 6 and 12 months. The data set includes all best ask and bid spot exchange rates for three major exchange rates: USD/EUR, USD/GBP and JPY/USD—hereafter EUR, GBP and JPY, respectively. It also includes ask and bid quotes for the exchange rate swaps for the four maturities as well as for euro-currency deposits for the four currencies involved.

An advantage of using deposit rates for interest rates is that an arbitrageur would know when and how much she will pay or receive. The use of deposits implies, however, that we limit the pool of potential arbitrageurs to those that have credit agreements, since deposits are on-balance sheet instruments. This limitation is not particularly severe in the present context since all major banks have such credit agreements established between themselves.⁹

⁸Ideally, in the case of JPY/USD we should have converted 1/10 of a pip in US dollars to JPY at the appropriate exchange rate at the end of each month—see Appendix A.B. On the other hand, we are probably deducting more than the average cost for each arbitrage deal involving three trades.

⁹Market participants without such credit agreement would instead use interest rate futures together with interest rate swaps. This will, however, incur uncertainty about exact amounts received at maturity.

For the spot exchange rates we have firm quotes from Reuters electronic brokerage system (D3000-2); these quotes are tradable as spot transactions can be carried out with a market order in the Reuters system. For swaps and euro-currency deposits only indicative ask and bid quotes were available to us through Reuters Monitor (i.e. Reuters 3000 Xtra). This is mainly because, both swaps and deposits are primarily traded bilaterally between interbank dealers, typically over telephone or Reuters D2000-1. Data from these sources is virtually impossible to obtain and has never been retrieved for empirical work in this context, to the best of our knowledge. Recently, an electronic-broker trading platform for swaps has been introduced, but this has yet to develop as the preferred platform.

However, in light of evidence for spot exchange rates in Danielsson and Payne (2002) and conversations with users of the Reuters trading system, one may say that spreads between indicative ask and bid quotes for swaps as well as for interest rates will not be smaller than those for corresponding firm ask and bid quotes.¹⁰ Thus, use of the indicative quotes would probably not lead us to exaggerate the number and size of arbitrage opportunities.¹¹ Actually, we may obtain results quite close to those implied by (unavailable) firm quotes for swaps and euro-currency deposits. This is because indicative quotes for swaps and deposits are also used for signaling in the dealer market, and hence regarded as a reliable indication of firm quotes in lack of other information sources. Essentially, because trading in swaps and euro-currency deposits only rarely occur on the Reuters electronic broker system, traders form their trading strategies on the posted indicative quotes, and the quotes on Reuters D2000-1 tend to be, therefore, very close to the quotes on Reuters 3000 Xtra. In contrast, indicative quotes for spot, on the other hand, are primarily meant as advertisement towards the non-bank customers and therefore not reliable indication of firm inter-dealer quotes. Thus, it is more important to have firm quotes for spot exchange rates than for exchange rate swaps and euro-currency deposits to obtain results close to those implied by firm quotes for all instruments.

Still, one drawback of using the indicative quotes is that they can become stale, at times, and thereby potentially signal spurious arbitrage opportunities. Usually, dealers keep real quotes up to date, but may fail to do so e.g. when market activity is particularly high, in which case indicative quotes will be centered on previous, rather than on current, firm quotes. If so, it will be possible for the indicative ask to be lower than the true ask or for the indicative bid to exceed the true bid, even while the indicative spread exceeds the

¹⁰Previous studies have also shown that the indicative mid-point is an accurate estimate of the firm mid-point; see e.g. Danielsson and Payne (2002).

¹¹In other words, if one cannot profit from arbitrage at the tight spread implied by firm prices, one definitely cannot profit at the worse prices in the wider spread provided by indicative quotes. The opposite is not true. In principle, therefore, there can be instances of arbitrage using firm prices that are not detected when using indicative prices.

true spread. Such cases could give rise to arbitrage opportunities in our empirical work when there are no opportunities in reality—i.e. false positives. We examine the robustness of our findings to this possibility by analyzing separately cases where an arbitrage opportunity arises because of a newly arrived spot quote in combination with existing swap quotes versus those that arise due to a newly arrived swap quote in combination with existing spot. The latter case is less likely to represent a stale swap quote.

In general, ask and bid quotes for an instrument (say the spot exchange rate) do not arrive contemporaneously with those for other instruments (e.g. euro-currency deposits for the currencies involved). In order to obtain continuous series of contemporaneous/synchronized (to the second) ask and bid quotes for different instruments, we merged all instruments according to date and time to the second into a file and then filled in missing ask and bid quotes for an instrument by using the latest quotes for that instrument. In order to severely limit the number of stale quotes, in our core empirical work we excluded weekends and days with unusually low or no trading activity (either due to a holiday or failure of the feed), which left us with quotes for 151 trading days.¹² In addition, we ignored quotes from hours with little trading and thus included only quotes that appeared during 07:00–18:00 GMT on the included days. In our robustness checks, we further limit the potential for stale quotes by imposing even more stringent constraints on how ‘fresh’ the quotes are and obtain largely the same results as in our core analysis.

Despite ignoring numerous observations to ensure calculations of arbitrage opportunities with as high a share of fresh quotes as possible, we are able to investigate a large number of data points (i.e. arbitrage returns), over 2 million in the case of EUR and around 2.5 million in the case of GBP. For JPY, however, about 0.8 million observations were obtained. The lower number of data points in the latter case can be explained on two grounds. First, our choice of trading hours allows us to cover trading in JPY taking place during the main European trading hours and partly the main US trading hours, at the expense of excluding the main Japanese trading hours. Second, the most active electronic market for trading JPY is the Electronic Broking System (EBS).¹³

4 Frequency, Size and Duration of Arbitrage Opportunities

In this section we report our key findings regarding the frequency, size and duration of arbitrage opportunities distinguishing between round-trip/covered arbitrage and one-way

¹²In addition to weekends, we left out the following days: April 2, 5–9, 12, May 3 and 31, June 17–18, August 10, 13, 24, and September 15, as these days were characterized by unusually low trading. Thus we were left with 151 days out of 231 days over the sample period February 13–September 30, 2004.

¹³ The EBS is also the main trading platform for EUR. Still, we have obtained a very large number of data points for EUR, although the largest number of observations is for GBP, for which Reuters is the main trading platform.

arbitrage. Our basic results are mere descriptions of the observations obtained by using the formulas described in Sections 2–3.

4.A Round-trip Arbitrage (CIP)

Table 4 presents results based on calculations of CIP arbitrage opportunities for the three exchange rates and four maturities examined. Results are given for both ask and bid sides—i.e. the outcomes of arbitrage both for the case when one borrows funds in the base currency to lend in the quoting currency and vice versa (these cases are referred to as “Ask” and “Bid” respectively, in the table). The table gives results for the case where all of the observations are used—Panel (a), “All deviations”—and for the case where only observations consistent with profitable deviations are considered—Panel (b), “Profitable deviations.” Starting from the case where all of the observations are used, we note that the number of observations increase with the maturity of contracts. This reflects that frequency of quote changes tends to be higher at higher maturities, especially for the swaps.

The table shows that the average return from CIP arbitrage is negative, in all of the cases—i.e. the figures in the column headed by “mean” are negative throughout the table. Also, the median return is very close to the mean return, indicating a fairly symmetric distribution. The negative mean values imply that, *on average*, CIP arbitrage is loss-making. Furthermore, the associated *t*-values suggest that the losses are statistically significant at conventional levels of significance.¹⁴ One would expect that arbitrage would eliminate any systematic negative or positive deviations from CIP and make CIP hold on average. One possible explanation for the negative mean of CIP deviations could be that market makers (quote providers) in the currency and deposit markets do not knowingly offer counterparts risk-free arbitrage opportunities and thus contribute to shift the returns towards negative values through their price offers. This would especially be the case if dealers, when pricing, say, the swap, worry about the fact that prices of other instruments, say deposits, may move in the next few seconds in a way to generate arbitrage. Accordingly, they may price more conservatively than CIP conditions imply in order to avoid arbitrage and be on the safe side. If prices are set in the deposit market in the same way, then equilibrium (average) prices will be consistent with a negative deviation from CIP rather than zero. An alternative explanation could be, that quote providers use the more stringent no-arbitrage conditions associated with OA and BA to set quotes (see Table 1). Accordingly, average returns in the case of CIP would be negative while those in the case of OA and BA will be zero. This latter explanation is consistent with the

¹⁴The *t*-values in the case of GBP are generally smaller in absolute terms than those for the other exchange rates, but still suggest significant losses in CIP arbitrage on average at the 5% level of significance.

results for the case of OA in BA in Tables 6 and 8, respectively, where we show that the LOP holds on average for both OA and BA conditions. Finally, by subtracting a generous amount of one tenth of a pip for other transaction cost we are biasing return downwards. Nevertheless, the negative average return from CIP arbitrage is not sufficient to prevent arbitrage in continuous time completely since the maximum point of the distribution of returns is not zero, which is the sufficient condition that is needed to prevent any arbitrage opportunity.¹⁵

The mean returns in Table 4 are period returns. It is therefore instructive to annualize them to make them more comparable across maturities. These calculations are given under the column headed “Ann. mean”, which illustrates how the (negative) returns are generally comparable across different maturities. In Table 4 we also document the pace of the market by “inter-quote time”, which is defined as the average time between two consecutive CIP deviations. Because at least one of the quotes involved in a CIP deviation formula must change in order to define a new CIP deviation, inter-quote time seems to be an appropriate aggregate indicator of the pace of FX and capital markets. The figures reported indicate that the pace of the market is very fast, especially at the higher maturities. New CIP deviations occur every 2-3 seconds on average for EUR and GBP, and every 6-7 seconds for JPY.¹⁶

Turning to the case where we consider only profitable CIP deviations, the column headed “Pa dev.” reports the number of profitable arbitrage opportunities out of the total number of data points available (“All dev.”) calculated for each of the exchange rates and maturities considered. Profitable deviations from CIP arbitrage are defined as the subset of CIP deviations with values in excess of 0.1 pip. The results suggest thousands of profitable arbitrage opportunities for all exchange rates, at most of the maturities. A round-trip arbitrage opportunity may on average arrive at least every hour when the number of profitable deviations (“Pa dev.”) are greater than 1661 ($= 151 \times 11$). As shares of the total number of data points considered, however, the profitable arbitrage opportunities are miniscule. The shares range from zero to 1.5% in the case of EUR, from 0.2% to 2.4% for GBP, and from 0.1% to 0.5% for JPY.¹⁷

¹⁵As noted in Section 2.A, loss on borrowing in, for example, the base currency to invest in quoting currency deposits does not necessarily imply a profitable arbitrage opportunity in the reverse direction.

¹⁶Our data set does not include all quotes, only all the *best* quotes. If all quotes were included—i.e. also those further out (worse prices) in the limit-order book that probably won’t be available for trading unless extremely large volumes are traded—inter-quote time would be much lower.

¹⁷Table 4 also suggests that there are fewer profitable arbitrage opportunities with lending dollar funds than when lending funds in euro, sterling and yen. This tendency is implied by the relatively higher share of profitable arbitrage opportunities on the ask sides relative to the bid sides in the case of EUR and GBP and on the bid side relative to the ask side in the case of JPY. In the latter case, USD is the base currency (f in the formula), while USD is the quoting currency (d in the formula) in the former cases—see Table 2—so that the finding of more arbitrage on the ask side of EUR and GBP and bid side of JPY might be

When examining the annualized mean return from profitable arbitrage deviations, we find that these returns range from a minimum of 2 pips in the case of EUR at the one-month bid side to a maximum of 15 pips for the JPY at the three-month ask. Also, the returns show no systematic pattern with maturity of the instruments involved in arbitrage.^{18,19} Finally, the average inter-quote time for profitable deviations ranges from less than 2 seconds to 15.6 seconds, except for one extreme case of 25 seconds for EUR at the one-month bid. In the latter case, the average inter-quote time is calculated only across 73 data points, which is the smallest number of arbitrage opportunities detected in Table 4.

Table 5 presents information about the duration of profitable CIP arbitrage opportunities. The table reports summary statistics of the durations of clusters (sequences) of profitable CIP deviations. A cluster is defined as consisting of at least two profitable CIP deviations in a row. The number of clusters, across exchange rates and maturities, ranges from a minimum of 8 to a maximum of 923. Notably, most clusters of profitable CIP deviations do not seem to last beyond a few minutes. Moreover, in most of the cases, average duration falls in the range from 30 seconds to less than 4 minutes. Median values of the durations are even lower than the corresponding average durations: they are generally less than 1 minute in the case of EUR; at most 1:43 minutes in the case of GBP; and at most 4:34 minutes in the case of Japan. It is worth noting that durations of clusters tend to decline, albeit non-monotonically, with the maturity of contracts. This seems to be consistent with the relatively high market pace (low inter-quote time) at higher maturities noted above.

Sample standard deviations of the durations reveal large variations in the duration of profitable CIP deviations, however. The standard deviations are quite different across the cases examined: they are mostly less than a few minutes, but exceptionally they can be

an artifact of the reverse quoting convention of dollar per euro and sterling vs. yen per dollar in Reuters system. There may be deeper microstructural reasons for this finding related to, for example, trade and execution issues, but we are not aware of any obvious reason why arbitrage opportunities occur more frequently when lending dollar funds over the sample analyzed. This is an intriguing issue which we leave for future research.

¹⁸The lack of relationship between size and maturity is in contrast with the conjecture that there may be a “maturity effect” such that the size of arbitrage profits increases with maturity. This conjecture was rationalized by Taylor (1989) on the basis of prudential credit limits that make arbitrage relatively more appealing at short maturities than at long maturities in a foreign exchange decentralized market where credit assessment is made cumbersome by lack of transparency. Of course, credit rating assessment is much easier within Reuters electronic system than in the pre-electronic, telephone-based brokerage systems studied by Taylor. For this reason, prudential credit limits may not provide a strong rationale for requiring larger returns for longer-maturity arbitrage activities in electronic systems such as Reuters. For a discussion of credit limits in decentralized and centralized, electronic markets, see Sarno and Taylor (2001, Ch. 2).

¹⁹In all cases, the median values of profits are comparable to the corresponding mean values, which also suggests fairly symmetric distributions of profits from round-trip arbitrage.

higher than 10 minutes. Often the relatively large standard deviations occur when there are relatively few observations, i.e. clusters. The first and third quantiles in the last two columns of Table 5 indicate that duration is not particularly high even at these quantiles of the distribution of durations, suggesting that the high standard deviations reported are potentially driven by relatively few outliers. They also explain the particularly long average duration of a few clusters of profitable CIP deviations.

Overall, we find a number of CIP arbitrage opportunities. However, these opportunities amount to small numbers when one compares them to the total number of observations examined. This is consistent with the widely held view that CIP holds tightly and that CIP violations occur rarely. The size of profitable CIP deviations is, however, economically appealing with periods returns (the annualized mean returns of profitable CIP arbitrage in Table 4) up to 15 pips. These are relatively large returns when compared with e.g. the typical size of spreads in the dealer markets for the major money and foreign exchange; usually around 2 pips. The size of the returns may seem small relative to the returns targeted by major players in the FX market, such as hedge funds, but they may not seem small if we take into account that they are riskless and require no own capital. The duration of profitable CIP deviations is relatively low but sufficiently high on average for a trader to exploit the arbitrage opportunities.

In order to exploit an arbitrage opportunity, however, a trader needs to undertake the three deals virtually simultaneously and as fast as possible.²⁰ Otherwise, there is a risk that prices of one or more instruments move such that an apparent arbitrage opportunity disappears before the trader has been able to seal all of the three deals—this may be termed ‘execution risk’. Reuters electronic trading system, which provides easy access to money and currency markets from one platform, allows a trader to undertake almost simultaneously several deals with counterparts. Alternatively, virtually simultaneous trading in the money markets and the swap markets can be accomplished through tight cooperation between money market dealers and swap market dealers which seems to exist in a typical dealing room.

We envision that a dealer observing an arbitrage opportunity would, given the non-negligible duration of profitable clusters, consider it worthwhile to inquire from her trading partners (including electronic broker for currency trading) about the relevant quotes that she would face, conditional on her (institution’s) credit rating and desired trade size. *Ex ante* the trader will not know for sure whether the quotes offered to her will imply

²⁰If the three deals are conducted consecutively from a single platform, it may take above one minute; a typical deal usually takes 25 seconds on Reuters dealing system, see Reuters (p. 114, 1999). Hence, the consecutive deals will involve the risk that one has to stop short of completing all of the deals required for arbitrage owing to an unfavorable change in the price of instruments that remain to be traded.

profitable arbitrage or not. It is possible that she would be offered quotes that do not imply an arbitrage opportunity because of relatively poor credit rating or desire for trading a relatively larger size than recently transacted. In addition, she has to take into account that if she trades currency through the broker, prices can move in an unfavorable direction, especially if she wants to trade large volumes relative to the standard or minimum size. If arbitrage remains profitable at the offered quotes while making allowance for sufficiently large currency price movements, and she is able to seal the required deals at those prices, the resulting profit will be risk-free. If these two conditions are not met, the trader will be at liberty to decline trading at the provided quotes, and hence will not suffer any loss.²¹ Danielsson and Payne (2002) observed that the likely price impact by “walking the book” in the major currency markets can be 1-2 pips when traded volumes increase from the minimum size of 10 million USD to, say, 30 million USD. Thus, a price impact in the currency markets is unlikely to remove a typical size of arbitrage profit, unless one makes an attempt to trade very large volumes.

Our findings suggest that arbitrage opportunities arise frequently enough and are sufficiently profitable to provide agents incentives to watch the markets and collect and analyze prices to the end of discovering and exploiting arbitrage opportunities. Moreover, due to their persistence, one may not need to continuously collect or inquire about prices that oneself would face, but only upon noting an arbitrage opportunity at prices that has been indicated to market participants in general or have recently been transacted at by other market participants.

Furthermore, our evidence suggests that the (average) size of profits is sufficiently large to reduce the possibility that a profitable arbitrage opportunity, at e.g., minimum trading size of 10 million USD at Reuters, will disappear just because one wants to trade a somewhat higher amount. Thus, it would be worthwhile for a trader to inquire about quotes even for trading higher volumes than those recently traded, for the purpose of undertaking arbitrage if considered profitable at the provided quotes.

4.B One-way Arbitrage: Owner Arbitrage (OA) and Borrower Arbitrage (BA)

We now report results for both forms of one-way arbitrage, i.e. for both OA and BA. The results are presented in the same format as for the CIP arbitrage reported in Tables 4-5.

Table 6 reports characteristics of OA opportunities for the three exchange rates and four maturities considered. OA calculations deliver period returns that are generally

²¹It should be borne in mind that a trader is free to only act on quotes that would benefit her, and is not obliged to act on every provided quote. Also, a trader does not have to make all sell and buy orders without knowing beforehand at what prices the orders will be executed at.

negative, but always insignificantly different from zero, on the basis of t -values. This is in contrast with the CIP deviations, and suggests that in the case of lending services, the LOP holds *on average*. The median returns are mostly close to the mean values, as in the case of CIP deviations. However, in contrast to the case of CIP arbitrage, the results indicate the presence of a large number of profitable OA opportunities in most cases. In particular, the frequency of profitable OA, calculated as shares of profitable OA opportunities out of the total number of deviations available, is in the range from about 15% to 48% in the case of EUR, 12% to 46% for GBP, and from about 11% to over 64% in the case of JPY.

As in the case of CIP arbitrage, there does not appear to be any systematic pattern linking returns to maturity. The annualized riskless mean returns from OA range from a minimum of 2 pips to a maximum of 6 pips across the three exchange rates examined. We also note that in contrast to the distributions of all deviations for OA, distributions of profitable OA opportunities seem to be slightly skewed to the left as the median values of the gains are generally lower than the corresponding mean values. The average inter-quote time is in the range from about 2 seconds to 9 seconds, confirming the fast pace of FX and capital markets. It also supports the tendency of longer-maturity markets to display faster activity levels than shorter-maturity markets.

Table 7 presents information about durations of profitable OA opportunities. The summary statistics of durations are quite similar to the case of CIP violations. In particular, means of the cluster durations are mostly less than about 5 minutes, and seem to decline with the maturity of the contracts. The standard deviations of cluster durations display more variation across the different cases, in comparison with the case of CIP, while median durations are lower than the corresponding mean durations.

Let us now turn to BA opportunities, analyzed in Table 8. On average, BA returns are generally negatively signed and always insignificantly different from zero. This suggests that the LOP holds on average also in the case of financing services. The frequency of BA opportunities is largely comparable to that of OA opportunities, with the corresponding shares ranging from about 8% to 50% in the case of EUR, 13% to 50% for GBP, and from about 11% to 68% in the case of JPY. The average sizes of gains from BA opportunities are also comparable to those from OA in Table 6. Moreover, the distributions of the gains from BA is slightly skewed to the left, as the median values are generally lower than the corresponding mean values. Furthermore, the annualized riskless mean returns from BA also range from a minimum of 2 pips to a maximum of 6 pips across the three exchange rates examined. In addition, the inter-quote times are similar to those in the case of OA, in the range from about 2 seconds to about 9 seconds. Finally, Table 9 shows that the

summary statistics of the cluster durations are similar to the case of CIP and especially to the results for OA.

Overall, we find a large number and high share of opportunities for one-way arbitrage—both in the form of OA and BA opportunities. The relative higher shares of OA and BA relative to CIP arbitrage are consistent with the implied relationship between CIP and OA forms of arbitrage discussed in Section 2. They are also consistent with our impression based on conversations with several FX dealers, who were of the view that OA and BA opportunities do arise much more frequently than CIP violations. Dealers may consider OA form of arbitrage whenever they receive funds to allocate, while they may consider the BA form of arbitrage when looking for funding opportunities. Another reason for the relatively higher numbers and shares of one-way arbitrage opportunities than those of round-trip opportunities could be due to the fact that one-way arbitrage requires excess supply (desire to lend) or demand of (desire to borrow) funds, while round-trip arbitrage requires no own funds or borrowing needs. Then, if the price elasticity of supply and/or demand is finite, one-way arbitrage opportunities, i.e. deviations from the LOP, are quite likely to appear regularly.

Furthermore, the relatively high shares of OA and BA deviations could also be due to institutional reasons, relating to the structure and trading practices of big corporations. Specifically, if we start from the outer circle of a general equilibrium model where corporations provide credit and funding to economic agents in the economy, each corporation will seek to minimize the cost of borrowing and to maximize the interest earned on surplus funds. Corporations will transact at different prices and a corporation’s credit spread changes continuously as its securities trade. One-way arbitrage would have to be part of this circle for it to matter to these corporations, but if indeed agents transact at different rates it is unlikely to have one equality (parity) holding for all agents. Turning to the inner circle of FX and money markets, the only way to justify the existence of one-way arbitrage opportunities is if the FX market is actually used by these agents to lower the cost of borrowing or earn extra yield. Yet there is little evidence of their treasury departments actively participating in these markets, which are largely dominated by individual FX and money market desks, whose capital is turned over fast with the objective of not having excess funding or borrowing needs and is used to support active trading, or ‘scalping’ the higher spreads faced by the broader corporate institutional market relative to those in the inter-dealer market. This interpretation of the economics of one-way arbitrage suggests that it is plausible that most dealers do not actively engage in the funding cost minimization or yield maximization (i.e. one-way arbitrage) and this is perhaps one reason of why we observe relatively high shares of OA and BA deviations.

The two one-way arbitrage opportunities display similar properties in terms of both size and duration. The size is less economically appealing than CIP deviations—especially considering that OA and BA are not pure arbitrage opportunities. Their duration is relatively low, consistent with the notion that these opportunities are short-lived, and tends to decline with the maturity of contracts, presumably because the market pace is higher for longer-maturity contracts. However, the relatively shorter duration of one-way arbitrage opportunities do not, *per se*, make them less exploitable than round-trip arbitrage opportunities. Opposed to the case of CIP, one-way arbitrage only requires two virtually simultaneous deals in order to be risk-free. Hence, a dealer able to undertake the two deals simultaneously, i.e. in a position to complete the required transactions at the prices implying OA or BA, will be able to obtain higher return or lower costs, respectively, through one-way arbitrage.

However, while the substantial shares of OA and BA opportunities is indicative of mispricing, i.e. non-zero inter-market price differentials and hence violation of the LOP, the existence of these opportunities has a different economic meaning relative to the less frequent CIP violations detected earlier. The CIP violations appear consistent with the Grossman-Stiglitz view of financial markets, whereby very short-term arbitrage opportunities arise, inviting traders to exploit them, and thereby be quickly eliminated. In the context of pure arbitrage, CIP violations are part of the interaction among traders setting prices in different but related markets which leads to correct equilibrium prices. In the context of OA and BA opportunities, these are LOP violations which have less important consequences for equilibrium prices since they do not imply pure arbitrage and hence would probably not receive a lot of attention from FX traders or commercial corporations.²²

In this sense, the large share of one-way arbitrage opportunities recorded does not necessarily imply that dealers are ill-informed and/or irrational. In fact, the apparent “mispricing” could reflect deliberate actions of well-informed rational dealers who actively manage their inventories of financial assets through their price setting, as is well known from the microstructure literature on quote-shading (e.g. Garman, 1976; O’Hara, 1995). Accordingly, dealers acting as market makers may deliberately provide relatively lower ask quotes and bid quotes if they want to reduce their inventories, or provide relatively higher ask and bid quotes if they want to increase their inventories. This explanation

²²It is also not clear what category of agents would be particularly inclined to search regularly for lower borrowing rates or higher lending rates. In general, unless the inter-market price differentials were very large (say over 5 pips) and the capital involved was extremely large the saving gained from borrowing more cheaply or from lending at the highest rate available is likely to be economically small. For example, it seems unlikely that an institutional player would care much about a price differential of 2 pips (saving 2000 dollars) on a capital of 10 million invested in a 12-month commercial paper issue. However, a CIP arbitrageur would exploit a mispricing of this size, but of course in the case of CIP arbitrage there are no capital requirements or lending needs.

of the apparent mispricing is not implausible given that a substantial share of trading in financial markets, and especially in the inter-dealer markets, is aimed at controlling inventories (e.g. Lyons, 2001; Bjønnes and Rime, 2005). A further exploration of this and alternative explanations of the numerous one-way arbitrage opportunities observed in our data set is left to future research.²³

5 Robustness and Further Empirical Analysis

In this section we undertake some robustness checks on the results reported in the previous section and further empirical work designed to shed light on the properties of arbitrage violations and their relation to the pace of the market. We start by investigating two important issues in this context, namely we assess how genuine the arbitrage opportunities detected are—i.e. we examine to what extent the noted arbitrage opportunities may be caused by stale quotes—and the possibility that the LOP in the FX market is guaranteed by prices being set directly from the no-arbitrage conditions for the cases of CIP, OA or BA.

5.A Are Arbitrage Opportunities Genuine or Due to Stale Quotes?

As described in our data section, we restricted our core analysis to the most active periods of market activity in order to limit the possibility of using stale quotes, that is quotes that may not be actually tradable even if they appear on the Reuters system. This was achieved by excluding weekends and days with unusually low or no trading activity (either due to a holiday or failure of the feed), and by including only quotes during the highest activity part of the trading day, namely 07:00–18:00 GMT.

In this sub-section we address the sensitivity of our results by further restricting the sample to quotes that may be considered particularly “fresh”. To this end, we amend the data set used until now as follows: we consider a quote of an instrument as “stale” if it occurs in an inactive state of the market, more precisely that the quote has not changed within the last two minutes, and will not change within the next two minutes.

²³An alternative explanation of the one-way arbitrage opportunities recorded here is the following: It starts with the observation that in cross-currency swap markets there is a known phenomenon of “points” being charged on one side of the swap due to the supply and demand for flows in different currencies. In essence, the exchanges are deliberately carried out at a profit to one side to compensate for taking the heavily supplied side of the trade (e.g. when the vast majority of dealers want to take the same long-short position in a currency pair). Given that the deposit rates are not adjusted by this price of aggregate imbalance, one may well find LOP violations that are spurious as they do not fully reflect the risk of the position of dealers taking the opposite side to what is demanded by the vast majority of agents. While these “points” are large for long-maturity swaps, it is not clear how large they may be for short-term transactions. However, their existence suggests that swap traders and arbitrage desks—as opposed to spot and money market dealers—may be the main players engaging in CIP, weakening the case for the relevance of one-way arbitrage opportunities.

When calculating deviations from CIP, OA and BA we now require that the quotes of *all* instruments involved in a arbitrage opportunity are fresh, i.e. they are not stale according to the above definition. All deviations which did not meet this criterion were excluded from the sample. Note that this a very stringent condition which should ensure we select data at very active trading times because, while quotes for spot exchange rates and for swap rates change at very fast pace, quotes for euro-currency deposits change at a significantly slower pace. Consequently, the results, reported in Table 10, indicate that this screening of the data reduces drastically the number of observations analyzed, especially for JPY. Also, the number of profitable deviations decreases substantially.

However, the frequency of occurrence of arbitrage—calculated as the share of profitable arbitrage violations out of the total number of deviations based on the particularly fresh quotes—remains fairly similar to the frequencies reported for CIP in Table 4, for OA in Table 6, and for BA in Table 8. Specifically, we find a low share for CIP violations, ranging from zero to about 3%, and high shares for both OA and BA, ranging from zero to 83%, and from zero to 76%, respectively. We also note that for some exchange rates and maturities the frequency is lower in this selective data set, while in some cases the frequency is higher relative to the baseline data set used in Section 4.

The filter above ensures that all quotes are fresh. In order to also control for the possibility that one market, most likely the spot market, moves fast and indicative quotes become “off-market”, we also repeat the exercise by adding the condition that the two consecutive spot rate quotes are identical. The results, available upon request, are similar to the results reported above.

Overall, these findings corroborate the results in Section 4 and add credibility to the view that the CIP arbitrage and LOP violations reported here are genuine opportunities which traders may have been able to exploit at the time of their occurrence.²⁴

5.B Is Any of the Assets Involved in Arbitrage Priced Using No-arbitrage Conditions?

Given our findings that arbitrage violations exist, at least one of the assets involved in FX arbitrage must sometimes be mispriced to an extent that is sufficient to generate arbitrage opportunities. With tick data on all four assets at our disposal, we can assess the mispricing in each asset to shed further light on how arbitrage arises.

Anecdotal evidence suggests that forward contracts may be priced according to the CIP condition, or even more stringent conditions (e.g. OA or BA). If this is the case, then

²⁴We also carried out the exercise in this sub-section by using fresh quotes defined on the basis of a 1-minute (rather than 2-minute) change in price, and found results that are qualitatively identical to the ones reported in Table 10. These results are not reported but remain available from the authors upon request.

price setting would be carried out in such a way as to prevent arbitrage opportunities from arising. We are in a position to test this conjecture in continuous time for the first time in this literature. More generally, we can test whether any of the four assets involved in arbitrage—spot, swap, domestic and foreign interest rates—is priced according to no-arbitrage conditions in continuous time. Given our data on CIP, OA and BA opportunities, we calculated the number of times an arbitrage (round-trip or one-way) opportunity was present at the same second when only the market swap quote was fresh (just posted), whilst the quotes entering the derived swap points were predetermined (or did not change). We then carry out the same exercise for the case when, in turn, only one of the spot exchange rate, the foreign interest rate and domestic rate are fresh quotes. We would expect that if an instrument were priced using, e.g., the CIP formula, the CIP condition should be valid at least whenever that instrument is priced, i.e. whenever the quote for that instrument changes.

Our results, reported in Table 11 for the case of CIP arbitrage, do not support this conjecture. The table shows that the profitable deviations owing exclusively to a new quote in one of the four instruments involved in arbitrage when the other instruments do not change (even if they are quoted) as a share of all such possible cases is comparable across instruments and also comparable to the corresponding shares reported in Table 4. These results do not support the view that forward rates—or spot rates or else deposit rates—are systematically set such that they ensure the validity of the CIP formulas using all available information. Thus, apparently, either the practice of using the no-arbitrage conditions to set prices is not feasible at this high frequency, and/or the providers of quotes do not update the formulas with all available information when offering quotes.

We also carried out this exercise for each of OA and BA (given in Tables 12 and 13 respectively), recording similar results to the ones reported in Table 11 for CIP. Overall, this evidence indicates that none of the asset prices is systematically set using no-arbitrage conditions at tick frequency, and that each of them is partly responsible for the mispricing leading to the CIP violations and one-way arbitrage opportunities reported here.

5.C Day of the Week and Hours of the Day Effects

In an attempt to understand the degree of predictability of CIP violations and one-way arbitrage opportunities and the existence of hours-of-the-day and day-of-the-week effects, we investigated whether CIP, OA and BA violations occur during a specific time of the day when, for example, trading might be particularly low—e.g. at the very beginning of the trading day examined. We also examined whether they occur more frequently in some days of the week than others. The results are reported in Tables 14 and 15 for day-of-

the-week and hours-of-the-day effects, respectively. The tables only report the results for CIP arbitrage to conserve space as the results for OA and BA are qualitatively identical.

Upon examining the share of CIP arbitrage across days (Table 14), we find that arbitrage violations occur throughout the week days (Monday to Friday) considered here, suggesting the existence of no systematic patterns of CIP violations during the week. Specifically, it is never the case that CIP arbitrage is more frequent across maturities or exchange rates in any given day relative to another. Similarly, the results in Table 15, reporting the share of CIP arbitrage during the trading day from 07:00 to 18:00 GMT, indicate no clear-cut evidence of clustering at a particular time of the day. However, in general, CIP opportunities appear to be somewhat more frequent early in the trading day and around the middle of the day, presumably because market pace is lower and quotes are updated less frequently at those times.

Overall, these results lend support to the findings in Section 4 that these opportunities are genuine short-lived arbitrage opportunities that may occur at any time.

5.D The Role of Market Pace and Volatility

Given the high frequency of the data employed in this study, it is difficult to design a comprehensive empirical analysis of the economic conditions under which arbitrage arises as most economic and financial variables are not available at this frequency. However, in this sub-section, we provide some illustrative evidence on whether frequency, size and duration of profitable arbitrage opportunities vary with the pace of the market and with market volatility. We undertake this investigation for all forms of arbitrage considered above to examine the generality of the findings. To this end, we estimate simple linear cross-sectional regression models with measures of frequency (share), size and duration of profitable CIP violations and one-way arbitrage (OA and BA) opportunities as dependent variables, regressed on an intercept, inter-quote time and a proxy for market volatility as the explanatory variables. That is, we estimate regression models of the following form:

$$y_i = \alpha_y + \beta_y IQ_{y,i} + \gamma_y DiffImplVol_{y,i} + \varepsilon_{y,i} \quad (9)$$

where $y = Share, Size, \text{ or } Duration$ of deviations from no-arbitrage conditions; IQ denotes inter-quote time; $DiffImplVol_{y,i}$ is the difference between maximum and minimum implied volatility and is a measure of the degree of uncertainty (variability) in volatility; and ε_y is an error term. Subscript i indicates an observation number; $i = 1, 2, 3, \dots, N_y$. The Greek letters represent time-invariant parameters.

The models are estimated by ordinary least squares (OLS) for each of the currency pairs examined. Accordingly, values for y , IQ and $DiffImplVol_{y,i}$ as well as the total numbers

of observations (N) depend on the form of arbitrage and the currency pair analyzed. We obtained observations for y , IQ and $DiffImplVol_{y,i}$ and stacked these in corresponding columns in accordance with both the arbitrage direction (i.e. stacking together ask and bid sides) and the maturity of the instruments involved.²⁵ Thus, the total number of observations N_y becomes equal to the sum of the total number of observations associated with the different maturities for each y_i examined.

The variables are defined more precisely as follows. The y -variable *Share* is defined, for a given form of arbitrage and currency pair, as the share of profitable deviations out of the total number of deviations from the corresponding no-arbitrage condition that occur in a business hour over the sample period. In this case, N_y can potentially be 13,288, which is the product of the 2 potential arbitrage directions (ask and bid); the 4 maturities considered; the 11 business hours (between 07.00–18.00 GMT); and 151 working days included in the sample. However, profitable arbitrage opportunities neither occur every hour in our sample nor in both directions. Thus, N_y is expected to be much lower than 13,288, especially in the case of CIP arbitrage. Each observation of IQ in the regressions for a specific currency pair would be equal to the average time between all of the (profitable and non-profitable) deviations used when calculating the corresponding observations for that frequency. Similarly, each observation of $DiffImplVol_{y,i}$ for frequency for a specific currency pair would be equal to the difference between the maximum implied volatility and the minimum implied volatility for all of the (profitable and non-profitable) deviations used when calculating the corresponding observations for that frequency.

The y -variable *Size* measures the average return of profitable deviations in a profitable cluster, while *Duration* refers to the time a profitable cluster lasts. The IQ variable in the regressions for *Size* and *Duration* refers to the average time between the row of profitable deviations constituting a profitable cluster, whereas the variable $DiffImplVol_{y,i}$ refers to the corresponding difference between the maximum implied volatility and the minimum implied volatility within each cluster. For a given form of arbitrage, the total number of observations used in a regression for size or duration for a currency pair would be equal to all profitable clusters for that currency pair.

The results from estimating regression (9) for frequency, size and duration, for all three currency pairs and no-arbitrage conditions, are given in Table 16. The results suggest that these characteristics of CIP arbitrage violations and one-way arbitrage opportunities tend to vary with the pace of the market, as proxied by the inter-quote time, and with the variability of volatility, as proxied by the difference between the maximum implied volatility and the minimum implied volatility. In particular, frequency, size and duration

²⁵Alternatively, we could have formulated separate models for the ask and the bid sides and for each of the four maturities examined.

are positively related both to IQ , i.e. negatively related to the market pace, and to $DiffImplVol_{y,i}$, i.e. positively related to variability of volatility. This suggests that when markets are particularly active, as described by a high number of new quotes per unit of time, and when the degree of uncertainty is relatively more stable, we should observe fewer arbitrage opportunities, smaller arbitrage profits, and more short-lived arbitrage.

For frequency and size there are a few instances where IQ or $DiffImplVol_{y,i}$ do not enter the regression with a statistically significant coefficient. However, the results are particularly clear-cut in the case of duration—in terms of both obtaining correctly signed and statistically significant positive coefficients. This suggests that high market pace and stable volatility are intimately related to arbitrage opportunities that are more short-lived, as one would expect.

6 Conclusions

Finance theory postulates that in well-functioning markets no-arbitrage conditions hold in continuous time, not just on average. This paper provides evidence that short-lived pure arbitrage opportunities arise in the major FX and capital markets in the form of violations of the CIP condition. The size of CIP arbitrage opportunities is economically significant for the three exchange rates examined and across different maturities of the instruments involved in arbitrage. The duration of arbitrage opportunities is, on average, high enough to allow agents to exploit deviations from the CIP condition. However, duration is low enough to suggest that markets exploit arbitrage opportunities rapidly. The high speed of arbitrage recorded in this paper can explain why such opportunities have gone undetected in much previous research using data at lower frequency.

We could detect the existence and measure the duration of a number of short-lived arbitrage opportunities only by using a unique data set at tick frequency for quotes of comparable domestic and foreign interest rates and spot and forward exchange rates. These features of the data set have proven essential to establish whether deviations from no-arbitrage conditions actually represented a profitable opportunity to agents at a given time or not, and to shed light on the time the market requires to restore no-arbitrage prices in an electronic trading platform such as Reuters. In turn, this emphasizes why studies of arbitrage require the analysis of tick, carefully matched data on the assets involved in arbitrage with meticulous attention to the finest institutional details.

We also report evidence that violations of one-way arbitrage opportunities arise frequently. These are violations of the LOP in the sense that they represent violations of the condition that the domestic lending (borrowing) interest rate should be the same as the foreign lending (borrowing) interest rate when the latter is adjusted to fully hedge

for exchange rate risk. The concept of one-way arbitrage is fundamentally different from CIP arbitrage and, more generally, the notion of pure arbitrage, i.e. instantaneous riskless profit net of all costs. This is because violations of these conditions do not imply riskless profits, but they do imply that the same need–desire to lend (owner arbitrage) or to borrow (borrower arbitrage)—can be satisfied at two different prices at a point in time. However, the evidence on the existence of one-way arbitrage opportunities is consistent with the general evidence of mispricing that characterizes these markets in continuous time and in contrast with the simplest finance textbook notion that dealers price assets according to the LOP to generate equilibrium prices instantaneously.

We find it comforting that the observed short-lived arbitrage opportunities provide evidence in support of the resolution proposed for the Grossman-Stiglitz ‘arbitrage paradox’. If arbitrage was never observed, market participants would not have sufficient incentives to watch the market, in which case arbitrage opportunities could arise. In turn, very short-term arbitrage opportunities invite traders to exploit them and hence will be quickly eliminated. While this view of arbitrage is appealing, previous empirical studies have been unable to detect such short-term arbitrage opportunities in a variety of financial markets, and this paper explains why that can occur when temporally aggregated data are used. To reiterate, arbitrage is indeed very short-lived, and requires turning on the microscope on high-quality tick data to be detected. This is because arbitrage-free prices are restored quickly, generally consistent with the notion of market efficiency.

A Appendix: Details on Calculations and Transaction Costs

A.A Calculating Days to Maturity (D)

We adjust interest rates, which are quoted in per cent per annum, by $D/360$ or $D/365$ to obtain interest rates for a period of less than a year. By convention, 365 refers to the total number of days in a year for a Commonwealth country, while 360 refers to the number of total days for other countries. D is the actual number of business days between the (spot) value date and the maturity date, which is generally the same date as the value date but in a different month.

Exceptionally, if the maturity date is a holiday in the home country of a security, the maturity date becomes the first business day after that holiday. If the value date is the last business day in a month, the maturity date will also be the last business day but in a different month. This is commonly referred to as the “end-of-month end-of-month rule.” For swap contracts, the value date and the maturity date must not be a holiday in the US and in the home countries of the quoting and the base currencies. We took holidays, i.e. days that are not settlement dates, for the different currencies from Bloomberg to account for this convention.

For almost all securities the value date falls on the second business day after the day of trading. The exception is the Eurosterling interest rate where the value date is the same as the trade date. Consequently, the maturity date of a sterling security that is traded on the same date as, e.g., a dollar security would, generally differ by two days. In order to ensure that both securities mature on the same day, dealers borrow or lend a sterling security forward with maturity on the value date of the other currency. Such deals are made through direct contact between dealers and, hence, do not generate transaction costs payable to Reuters.

A.B Transaction Costs: Brokerage Fees and Settlement Costs

There are two types of variable transaction costs associated with trading in the FX market, in addition to those captured by ask-bid spreads: brokerage fees and settlement costs. In our case, the brokerage fees refer to the costs of trading swap contracts through the Reuters electronic broking system, Reuters Dealing 3000. At present, the Reuters system does not allow for trading of deposits in the security markets. Such trading is conducted via direct contact between dealers or through voice brokers. The variable broker costs of trading in deposits may therefore be assumed to be zero. Settlement costs, however, are incurred on trades of both swap contracts and deposits.

The brokerage fee is paid by the initiator of a trade (aggressor) at the end of a month in the Reuters trading system for swaps. Such fees increase with the maturity of a traded swap contract, but are inversely related to the total volume traded by the aggressor in a month. Table A.1 presents a recent fee schedule for Reuters dealing system, where we report deal fees charged when dealing swaps through Reuters Dealing 3000. When a volume band has been reached, the (lower) deal fee per million (mill) USD in the subsequent band is applied to the total volume. It appears that a small trader with a total trade volume of 1 billion (bn) USD or less incurs a fee of at most 10 USD for a trade of 1 mill USD at maturities of one month to one year (inclusive). If one trades more than 5 bn a month in this maturity range, the fee falls to 5 USD for a trade of 1 mill USD.

An arbitrage deal using a currency swap leads to three transactions, one in the FX market and two in the security markets, and thus for a total of 9 ($= 3 \times 3$) notices. Hence, the total (variable) settlement costs vary in the range of $1.26 - 2.52$ ($= 3 \times 0.42 - 3 \times 0.84$) USD. In extraordinary situations, a trade may require more than three notices and, therefore, entail higher costs.

Overall, even the total of variable transaction costs (brokerage fees and settlement costs) per unit of a base currency becomes negligible. For example, the sum of brokerage fee and settlement costs of a minimum-size swap of 10 mill USD of, e.g., maturity one month to one year (inclusive), would at most be $(10 \times 10 + 0.84) = 100.84$ USD, i.e. 10.084 per 1 mill USD or about 1/10 of a pip per USD. If we add the SWIFT costs associated with lending and borrowing, the total cost associated with an arbitrage deal involving a minimum-size swap, would still be about 1/10 of a pip, or more precisely $(100.84 + 2 \times 0.84) = 10.252$ per 1 mill USD or 0.1025 of a pip per USD.

are invoiced in euros. Where fees are denominated in another currency, they are converted to dollars or euros at the market spot selling rate at 15.00 Belgian time at the end of the period for which the invoice is issued.

References

- Aliber, R.Z., Chowdry, B. and Yan, S. (2003), "Some Evidence that a Tobin Tax on Foreign Exchange Transactions May Increase Volatility," *European Finance Review*, 7, 481–510.
- Branson, W.H. (1969), "The Minimum Covered Interest Differential Needed for International Arbitrage Activity," *Journal of Political Economy*, 77, 1028–1035.
- Bjønnes, G.H. and Rime, D. (2005), "Dealer Behavior and Trading Systems in the Foreign Exchange Market," *Journal of Financial Economics*, 75, 571–605.
- Callier, P. (1981), "One Way Arbitrage, Foreign Exchange and Security Markets: A Note," *Journal of Finance*, 36, 1177–1186.
- Clinton, K. (1988), "Transaction costs and Covered Interest Arbitrage: Theory and Evidence," *Journal of Political Economy*, 96, 358–370.
- Covrig, V. and Melvin, M. (2002), "Asymmetric Information and Price Discovery in the FX Market: Does Tokyo Know More About the Yen?", *Journal of Empirical Finance*, 9, 271–285.
- Deardorff, A. V. (1979), "One-way Arbitrage and its Implications for the Foreign Exchange Markets," *Journal of Political Economy*, 87, 351–364.
- Danielsson, J. and Payne, R. (2002), "Liquidity Determination in an Order Driven Market," Financial Markets Group, London School of Economics, mimeo.
- Evans, M.D.D. (2002), "FX Trading and Exchange Rate Dynamics," *Journal of Finance*, 107, 2405–2447.
- Evans, M.D.D. and Lyons, R.K. (2002), "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy*, 110, 170–180.
- Evans, M.D.D. and Lyons, R.K. (2005), "Meese–Rogoff Redux: Micro-based Exchange Rate Forecasting," *American Economic Review Papers and Proceedings*, 95, 405–414.
- Fletcher, D.J. and Taylor, L.W. (1996), "SWAP Covered Interest Parity in Long-date Capital Markets," *Review of Economics and Statistics*, 78, 530–538.
- Frenkel, J.A. (1973), "Elasticities and the Interest Parity Theory," *Journal of Political Economy*, 81, 741–747.
- Frenkel, J.A. and Levich, R.M. (1975), "Covered Interest Arbitrage: Unexploited Profits," *Journal of Political Economy*, 83, 325–338.
- Frenkel, J.A. and Levich, R.M. (1977), "Transaction Costs and Interest Arbitrage: Tranquil Versus Turbulent Periods," *Journal of Political Economy*, 85, 1209–1226.
- Garman, M. (1976), "Market Microstructure," *Journal of Financial Economics*, 3, 257–275.

- Grossman, S.J. and Stiglitz, J.E. (1976), "Information and Competitive Price Systems," *American Economic Review*, 66, 246–252.
- Grossman, S.J. and Stiglitz, J.E. (1980), "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393–407.
- Juhl, T., Miles, W. and Weidenmier, M. (2006), "Covered Interest Arbitrage: Then vs Now," NBER Working Paper No. W10961, *Economica*, 73, 341–352.
- Lyons, R.K. (1995), "Tests of Microstructural Hypothesis in the Foreign Exchange Market," *Journal of Financial Economics*, 39, 321–351.
- Lyons, R.K. (2001), *The Microstructure Approach to Exchange Rates*. Cambridge, USA: MIT Press.
- Lyons, R.K. and Moore, M. (2005), "An Information Approach to International Currencies," NBER Working Paper No. 11220, University of California, Berkeley and Queen's University, Belfast, mimeo.
- O'Hara, M. (1995), *Microstructure Theory*. Oxford, UK: Blackwell Publishers.
- Osler, C.L. (2000), "Support for Resistance: Technical Analysis and Intraday Exchange Rates," Federal Reserve Bank of New York *Economic Policy Review*, 6, 53–67.
- Osler, C.L. (2003), "Currency Orders and Exchange Rate Dynamics: Explaining the Success of Technical Analysis," *Journal of Finance*, 58, 1791–1820.
- Osler, C.L. (2005), "Stop-loss Orders and Price Cascades in Currency Markets," *Journal of International Money and Finance*, 24, 219–241.
- Payne, R. (2003), "Informed Trade in Spot Foreign Exchange Markets: An Empirical Investigation," *Journal of International Economics*, 61, 307–329.
- Reuters (1999), *An Introduction to Foreign Exchange & Money Markets*, Reuters Financial Training Series, Singapore and New York: Wiley.
- Rhee, S.G. and Chang, R.P. (1992), "Intra-day Arbitrage Opportunities in Foreign Exchange and Eurocurrency Markets," *Journal of Finance*, 47, 363–379.
- Sarno, L. and Taylor, M.P. (2001), *The Microstructure of the Foreign Exchange Market. A Selective Survey of the Literature*, Princeton Studies in International Economics series, No. 89, Princeton, NJ: Princeton University Press.
- Shleifer, A. and Vishny, R. (1997), "The Limits of Arbitrage," *Journal of Finance*, 52, 35–55.
- Taylor, M.P. (1987), "Covered Interest Parity: A High-Frequency, High-Quality Data Study," *Economica*, 54, 429–438.
- Taylor, M.P. (1989), "Covered Interest Arbitrage and Market Turbulence," *Economic Journal*, 99, 376–391.

