COLLATERAL CONSTRAINTS AND
MICRO-BUBBLES IN FOREIGN EXCHANGE

Efraim Berkovich[1]
Manhattanville College

ABSTRACT:

When a risk-averse market maker in the foreign exchange market has an inventory imbalance, he posts prices which deviate from fair value. Noise-traders provide inventory shocks to the market maker which arbitrageurs can, in expectation, profitably exploit. As arbitrageur positions grow, collateral constraints on arbitrageurs incentivize the market maker to post prices which deviate further from fair value—a momentum effect in prices. This bubble pops after arbitrageur inventory is resolved or value is revealed. High-frequency price data provide some evidence of these "micro-bubbles" occurring on an intra-day time-scale.

WORKING DRAFT.

---

1 INTRODUCTION

The foreign exchange (FX) currency market is by far the largest world market by USD volume[2]. Menkhoff and Taylor (2007) review the FX market literature on technical trading, that is, trading on the basis of public trade history, and find that excess profit opportunities appear to exist. Because of the size of the market and because the majority of trading is done by professionals, this apparent violation of market efficiency is puzzling.

Another puzzle relates to the disposition effect in behavioral finance—the tendency of traders to take many small gains and a few large losses. Studies have found that those traders most prone to the disposition effect tend to significantly underperform. It is not clear how these "anti-geniuses" could accomplish this feat in an efficient market.

This paper attempts to provide a micro-structural model of a price process in the FX market. The model predicts that prices in the market have a trend bias when leveraged arbitrageurs take on larger positions. Since this trend occurs endogenously, it is not possible for arbitrageurs in the market to take advantage of it. Arbitrageurs create the trend by trading and tying up their collateral at the exact time they would benefit from being able to trade more. New collateral cannot quickly enter given the time-frame, so the trend persists.

In the model, a monopolistic market maker sets price for a currency pair. Without new information about asset values, the market maker's risk-aversion to accumulating asset inventory drives his price setting behavior. Two types of traders trade with the market maker: "noise" traders who trade based on uncorrelated portfolio shocks and risk-averse arbitrageurs. No trader possesses private information, and information about asset value is publically revealed in an information event after a number of periods. Since the market maker adjusts posted prices to compensate for his inventory risk, prices fluctuate due to the random arrival of order flows from noise traders. As prices deviate from fair value, arbitrageurs can trade in the direction of fair value if the price deviation is sufficient to compensate the arbitrageur. Arbitrageurs can realize a profit when prices move in the direction of their trade. When prices move against an arbitrageur's position, the arbitrageur decides whether to increase, close, or hold the position. Limited collateral may prevent arbitrageurs from increasing or even continuing to hold a position.

The market maker anticipates arbitrageur collateral strain and adjusts price and his own inventory accordingly. Since the market maker sets the price, he has an incentive to post a price which lowers expected surplus for arbitrageurs. Thus, increased arbitrageur inventory and leverage cause the market maker to optimally post prices increasingly farther from fair value.

---

[2] The Bank of International Settlements estimated the total daily volume of FX transactions at $4 trillion in 2010 and $5.3 trillion in 2013.

Trading by collaterally constrained arbitrageurs endogenously generates momentum in price. This "micro-bubble" eventually dissipates when either the asset value is revealed or arbitrageur inventory is reduced. In principle, it is difficult to detect these micro-bubbles without data on trader leverage. However, posted prices from one FX dealer provide some statistical evidence in corroboration of the model. The presented evidence is by no means conclusive since an empirical test of the model requires data on individual trades and accounts which were not available. Nonetheless, some indications of predicted effects appear. At a timescale of about 10-30 minutes, prices respond with momentum to large past moves and revert for small moves. As sampling frequency decreases to about 4 hours, momentum gives way to reversion possibly indicating that the micro-bubble has deflated.

1.1 RELATED LITERATURE: THEORY

Noise trader effects on price have been well-studied across many literatures. In one early work, Kyle (1985) puts forward a model of a market with arriving noise-traders. A number of studies explain market price anomalies by positing various models where irrational agents have a persistent effect on prices. While the efficient market hypothesis contends that the existence of non-rational traders does not materially affect prices because rational arbitrageurs trade against them, many researchers have put forward models which describe markets where prices deviate from fair value because of the presence of non-rational traders. Only a few works are cited below as this important literature is too vast to be reviewed here. DeLong, Shleifer, Summers, and Waldmann (1990a, 1990b) show that when a sufficiently large number of non-rational traders trade with trend then so do rational arbitrageurs, thereby creating a price deviation. Furthermore, non-rational traders may not be driven out of the market; they earn higher nominal returns as they do not correctly price risk for themselves. Abreu and Brunnermeier (2002) present a model of price bubble formation: arbitrageurs are unable to synchronize their attacks on mispricing and therefore expose themselves to risk while they wait for other arbitrageurs to join the attack. This fact allows a price bubble to grow.

It is well known that collateral constraints affect asset pricing. Shleifer and Vishny (1997) consider collateral effects on arbitrageur trading and find reduced trading. Gromb and Vayanos (2002) examine collateral-constrained arbitrageurs in segmented markets and claim that under some circumstances arbitrageurs take on too much risk from a social welfare perspective, as they do not internalize the externality that possible liquidations have on market prices. Garleanu and Pedersen (2011) show that differing margin requirements directly affect prices of nearly identical assets in credit markets.

This paper contributes to the literature on collaterally-constrained arbitrage. A closely related work in that literature by Kondor (2009) offers a model of collaterally constrained arbitrage where arbitrageurs trade across multiple periods. Prices depend on arbitrageur collateral and its

allocation under uncertain future arbitrage opportunities. Prices change endogenously from arbitrageur activity much as in the proposed model. Price deviates further from fair value as time proceeds and arbitrageur capital is depleted. One major difference between Kondor (2009) and the present work is the monopolistic market maker who sets price strategically in response to arbitrageur inventory.

The market maker described in this paper derives from work on currency market microstructure which describes a two-tiered market. According to this approach, put forward in Lyons (1997) and expanded in Lyons (2001), risk-averse market makers absorb customer orders and then work to offload inventory risk into the inter-dealer market. Because it may take a number of trades before position risk is allocated to the parties willing to hold it, this "hot potato" effect explains the very large daily volume in the foreign exchange market. Evans and Lyons (2002) explain that order flow to market makers can be a good estimator of market prices. A number of other researchers, such as Osler, Mende, and Menkhoff (2011), examine the effect this risk-passing two-tiered market has on public information and price formation.

1.2   RELATED LITERATURE: EMPIRICAL AND BEHAVIORAL

Menkhoff and Taylor (2007) provide support for the proposed model's implication that the FX market is not weak-form efficient—that prices may be predicted from past public history. In their review of the literature on FX technical trading, Menkhoff and Taylor (2007) explain the wide use of technical trading amongst market participants, point to its possible (risk-adjusted) profitability, and hypothesize that its apparent success may be due to its ability to draw out non-fundamental price determinants. Menkhoff and Schmidt (2005) survey German fund managers (primarily of equities) and find widespread use of technical trading strategies at short timescales. They find those managers who favor momentum strategies to be less risk-averse while those who favor contrarian strategies to be overconfident and prone to the disposition effect. Although the examined market is not the FX market, these findings are in-line with implications of the proposed model where lower risk-aversion makes trend-trading profitable and where counter-trend trades imply a pattern of profit-taking consistent with the disposition effect.

The trading pattern implied by the model may explain the well-studied disposition effect in behavioral finance—that biased traders tend to close out gains and hold on to losses. The effect was described by Shefrin and Statman (1985) and further explored by many others including: Odean (1998) for U.S. retail stock traders, Feng and Seasholes (2005) for Chinese stock traders, and Locke and Mann (2005) for professional futures traders. These studies all find that traders who exhibit this apparent bias tend to underperform. The unanswered question is how that is possible if prices are efficient.

When arbitrageurs trade counter-trend in expectation of a high probability of small gains after prices mean-revert, they are willing to hold a losing position until forced out by collateral

constraints (or an information event which changes the expected mean). The resulting pattern of trade is consistent with observations of a disposition effect. Since the market maker adjusts price in the direction of arbitrageur inventory, arbitrageurs who trade this way may lose surplus to the market maker and underperform.

2 MODEL

Asset prices are often described as a stochastic diffusion process with jump events. In FX markets, price jumps may come from rapidly disseminated information events such as, for instance, central bank announcements. The diffusion "noise" part of the price process may arise from the arrival of random order flow generated by so-called noise traders.

Consider a multi-period market in a single asset. The asset's value is not publically known until the last period when the value is revealed via an information event. In each period, traders may trade with a single market maker (MM). The MM has access to an inter-dealer market which is not available to traders.

The timing in a period proceeds as follows:
  (1) the MM posts prices at which he is willing to buy and sell the asset,
  (2) traders see the prices,
  (3) traders buy or sell the asset in any amount, and
  (4) the MM trades any amount of the asset in the inter-dealer market.
The MM commits to absorbing all order flow from traders at the posted price for that period. Traders may only trade with the MM, so bilateral trades are not allowed. If the MM does not want to hold inventory risk, he may offload an inventory imbalance to a decentralized inter-dealer market, as described in Lyons (1997, 2001). In addition, the MM may change his posted price in a later period in order to reduce the price risk of additional inventory accumulation, taking nominal expected losses on customer orders which reduce his inventory imbalance.

There are two types of traders: (a) noise traders who buy and sell due to independent idiosyncratic shocks and (b) arbitrageurs who trade strategically. No agent, including the market maker, has private information about asset value. However, the distribution of the asset's value is common knowledge.

The following assumptions hold:

1. The MM posts a single price rather than a bid-ask spread. Assume the MM earns normal profit from his market-making activities and that the bid-ask spread is not material to the overall price movement.[3]

---

[3] Ignoring the effect of bid-ask bounce—that is, last price changing as traders alternate hitting the bid and ask prices.

2. The utility functions of the MM and arbitrageur—given by $U(w)$ and $u(w)$, respectively—are common knowledge with $U' > 0,\ u' > 0,\ U'' < 0$ and $u'' < 0$.

3. Noise traders trade regardless of posted price. Their aggregated demand in period t is given by an i.i.d. random variable x_t which has a probability density denoted by $g(x)$. Noise trader demand may be positive, indicating net buying, or negative, indicating net selling by traders. To simplify analysis, assume the density is symmetric around 0.

4. Although he cannot present different prices to different traders, the MM can identify the type of traders who transact with him.[4] Importantly, the MM knows account characteristics such as capital in the trader's account. Information on the aggregate level of arbitrageur positions is public.

5. When the MM has an inventory imbalance, he may trade in the inter-dealer market. He is able to reduce his excess inventory by up to ϕ units per period. (He may choose to transact less than ϕ in order to bring his inventory to zero.) The asset price in the inter-dealer market is the fair value price of the asset.

6. There are no interest charges/payments to any agent. It is straightforward to add interest rates to the analysis but at the cost of increased notation. Since interest amounts are likely to be very low for short time-frames, this simplification seems innocuous.

The value of the asset, v, has a distribution over $[0, \infty]$ with density $f(v)$ and is independent of order flow. Fair value—that is, the mean—is given by $\bar{v}$. The realized value is revealed at an information event at period T; those agents holding the asset after the information event consume the realized value.

Although the timing of some market information events is known (*e.g.,* central bank announcements), other information events may arrive randomly. For this analysis, uncertain timing of information events is not examined.

One can extend the model to include the presence of informed traders. When arbitrageurs and the MM have common knowledge about the probability of the arrival of an informed trader, model dynamics remain similar, although prices may be quite volatile as in Brock and Shleifer (2012). This model variation was not adequately explored, but in the limited analysis performed it does not appear to materially affect the qualitative results.

2.1 MARKET MAKER

---

[4] The MM is typically the brokerage for the traders and therefore knows identifying information about them. While in the real-world, the MM may not know a trader's type/motivation with certainty, currency dealers may develop a sense of the likelihood of the type of trader from various account characteristics.

The MM must trade with all comers at his posted prices. He earns an unmodelled market spread or other transaction benefit which allows him to offset utility losses from market making. As the risk-averse MM does not want to speculate on the final value of the asset, his risk is in holding inventory when the asset's value is revealed. Consequently, at times, the MM may want to trade in the inter-dealer market at fair value to reduce his inventory risk.

The MM's inventory at time t is given by N_t, cash (in the lead currency) by c_t, and history known to the MM at time t by h_t. His posted price in period t is given by p_t, and the amount he decides to trade in the inter-dealer market at the fair price is given by b_t.

2.2   A LONELY MARKET: NO ARBITRAGEURS

As a baseline case to examine the MM's problem, consider a market with only the MM and noise traders. Since history provides no additional information in a market with only noise traders, the MM's choice of price p_t and inter-dealer trade b_t depends only on the state of his inventory.

The MM's inventory evolves according to $N_{t+1} = N_t - x_t + b(N_t, x_t)$. Incoming demand changes inventory by $-x_t$ as the MM must take the opposite position of incoming orders. Since the order flow to the MM is a martingale, it is clear that the optimal choice for the risk-averse MM is to reduce inventory quickly to zero. Trading in the inter-dealer market provides a benefit to the MM because his risky inventory is converted into its fair price equivalent. Therefore, the optimal inter-dealer trade is given by $b(n, x) = sign(x - n)min[\phi, |x - n|]$.

Assuming the MM has inventory N_{T-1} at period *T-1* and cash (in the lead currency) of c_{T-1}, he sets the price as

$$p_{T-1}^* \in \underset{p}{argmax} \int_0^\infty f(v) \int g(x)U\big(v\{N_{T-1} - x + b(N_{T-1}, x)\} + c_{T-1} + xp$$
$$- \bar{v}b(N_{T-1}, x)\big)dxdv$$

Since *U* is strictly concave and p enters linearly, the problem has a unique solution. Due to concavity of *U*, price increases in inventory, $\frac{dp}{dN} > 0$. The MM optimally deviates from fair price in the direction of his current inventory—for instance, if inventory is positive, optimal price is higher than fair value. These deviations are optimal since the risk-averse MM is concerned with the chance of a large loss on his inventory. He, therefore, reduces the risk of a large loss by charging a price which reduces the size of any potential loss. This response comes at the cost of a smaller gain if his inventory position turns out to be more valuable than expected. Price change with respect to initial cash balance c depends on risk-aversion in relation to wealth; for instance, for constant absolute risk aversion, $\frac{dp}{dc} = 0$.

The MM optimally maximizes his inter-dealer trade for inventory reduction. Therefore, $\frac{dp}{d\phi} < 0$ when $N_{T-1} < 0$ and $\frac{dp}{d\phi} > 0$ when $N_{T-1} > 0$. Trading on the inter-dealer market would allow profitable arbitrage if the MM could pick which traders he deals with in his own market. Because he cannot, the MM's expected cash profit is zero for any price he posts. The effect of inter-dealer trading is to cushion the MM's risk by transforming inventory into its fair value. Thus, with larger ϕ, the MM optimally posts prices closer to fair value.

Denote by $p_{T-1}^*(N)$ the optimal price for period $T\text{-}1$ given inventory N. The problem for period $T\text{-}2$ is

$$\max_p \int_0^1 f(v) \iint g(x')g(x)U(v\{N_{T-2} - x' - x + b(N_{T-2}, x') + b(N_{T-1}, x)\} + c_{T-2}$$
$$+ x'p + xp_{T-1}^*(N_{T-1}) - \bar{v}\{b(N_{T-2}, x') + b(N_{T-1}, x)\})dx'dx\,dv$$

Continuing in this way, the MM determines the optimal price for any period t. Note that if $\phi = 0$, inventory at period T becomes the sum of $T\text{-}t+1$ i.i.d. random variables—a martingale.

2.2.1   EXAMPLE: CARA UTILITY

Suppose that $T\text{=}4$, $\phi = 0$, and the MM has CARA utility given by $U(w) = -\sigma^{-1}e^{-\sigma w}$. Suppose further that $x_t = \pm 1$ with equal probability and v is 0 or 1 with equal probability. Since the problem is symmetric, it is sufficient to examine prices conditional on $x_1 = 1$ (*i.e.,* demand to purchase 1 unit of the asset). Posted price in $t\text{=}1$ is ½, and the MM determines price $p_2(\text{-}1)$, that is, price in the second period when inventory is -1 (since the MM sold the asset). In period $t\text{=}3$, inventory can be either -2 (when $x_2 = 1$) or 0 (when $x_2 = -1$). For $p_3(-2)$, the maximization problem can be written as

$$\max_{p_3} U\left(-3 + p_3 + p_2 + \frac{1}{2}\right) + U\left(-1 - p_3 + p_2 + \frac{1}{2}\right) + U\left(p_3 + p_2 + \frac{1}{2}\right) + U\left(-p_3 + p_2 + \frac{1}{2}\right)$$

Using first-order conditions the solution is

$$p_3(-2) = \frac{1}{2\sigma} log\left(\frac{e^{3\sigma} + 1}{e^{\sigma} + 1}\right)$$

For $p_3(0)$, zero inventory leads to the unsurprising result

$$p_3(0) = \frac{1}{2\sigma} log\left(\frac{e^{\sigma} + 1}{e^{-\sigma} + 1}\right) = 1/2$$

Substituting the above prices into the maximization problem for $t\text{=}2$ solves for $p_2(-1)$. The solution by first-order conditions is

$$p_2(-1) = \frac{1}{2\sigma} log\left(\frac{e^{-\sigma(-3+p_3(-2))} + e^{-\sigma(-1-p_3(-2))} + e^{-\sigma(p_3(-2))} + e^{-\sigma(-p_3(-2))}}{2e^{\sigma/2} + 2e^{-\sigma/2}}\right)$$

2.2.2   PRICE SENSITIVITY TO TIME

For $\phi \geq sup(x_t)$, the MM trades his entire inventory to the inter-dealer market in every period and therefore posts a price of fair value always. For smaller values of ϕ, the MM's problem has

path dependence to his inventory accumulation. Because a larger ϕ shrinks inventory in the current period, N_t, and reduces expected inventory in the following period, N_{t+1}, the optimal price is closer to fair value for larger ϕ in all periods. The sensitivity of price to changes in inventory (and inventory reduction through inter-dealer trade) increases as time approaches the information event at time T, since the MM has less chance to reduce his inventory imbalance.

Proposition: *At any given price, holding cash equal, for $0 < \phi < sup(x_t)$, $\left|\frac{dp_t}{dN_t}\right|$ increases and $\left|\frac{dp_t}{d\phi}\right|$ decreases as $t \to T$.*

Proof. The MM's utility depends on $E[N_T]$, highest expected utility is when $N_T = 0$, and utility decreases in $|N_t|$. $Pr[N_T = 0|N_t]$ strictly increases as $N_t \to 0$. Since demand shocks are symmetric, inventory imbalance is reduced in expectation through inter-dealer trade, that is $|E[N_{t+1}]| < |E[N_t|N_t \neq 0]|$. Therefore, $Pr[N_T = 0|t < T]$, where the current period is denoted by t, increases as t decreases. Thus, at period t compared to period $t+1$ with the same level of cash and inventory, the marginal benefit to the MM of deviating price from $\bar{v}$ to offset inventory risk is smaller while the marginal cost remains the same (due to expected cash losses). Therefore, the MM optimally picks a price closer to $\bar{v}$ in period t as compared to period $t+1$. A similar argument shows that a higher ϕ causes the MM to optimally pick a price closer to $\bar{v}$ in period t as compared to period $t+1$.∎

As an aside, the argument above provides insight into a model where the timing of the information event is uncertain. Since the MM's expected utility depends on inventory he holds at the time of the information event, higher inventory at every period poses a risk since there is positive probability of the game ending next period.

2.3   ARBITRAGEURS

Arbitrageurs have no private information about the final value of the asset, but, unlike noise traders, may choose to trade depending on posted prices. Two assumptions are imposed on arbitrageurs:
  1.   Each individual arbitrageur is small relative to the market.
  2.   Arbitrageurs cannot synchronize their activities.
These two assumptions set the arbitrageurs as price-takers in a monopolistic market with the MM as monopolist.

Order flow from arbitrageurs is additive to noise trader order flow. Putting aside for the moment the effect arbitrageur order flow has on the MM's posted price, one can examine various arbitrageur trading strategies from the standpoint of a single arbitrageur whose order flow is very small. The full space of trading strategies is quite large, so the following analysis examines just two types of trades: momentum trend trades which trade away from fair value and counter-trend trades which trade in the direction of fair value. When an arbitrageur takes a position, he anticipates three possible outcomes:

(a) He will be able to close it voluntarily in a later period.

(b) He will wind up holding the position until the information event.

(c) He will be forced out of the position due to collateral constraints.

The setup is symmetrical, so it is sufficient to examine behavior for the top half of the price range. Use $u(q)$ to mean the level of utility when the arbitrageur has initial wealth plus profit q.

2.3.1 MOMENTUM TRADING

An arbitrageur who buys in period 0 may choose to sell in period 1 if price moves up. On the other hand, if price moves down, the arbitrageur faces a choice to close out the position for a cash loss or to hold the asset in hopes of closing the position later. Holding the asset may entail an even bigger loss. Even if the purchase price is close to fair value, this strategy may entail a large utility loss to a risk-averse arbitrageur if he winds up holding the asset until period T.

Consider a purchase of one unit of the asset in period 0 at price p_0. Given public history h_{T-2} at the end of period T-2, the expected gain of holding the position and possibly selling for a gain in $t=T$-1 is

$$D^{sell}_{T-1}(h_{T-2}) = \Pr[p_{T-1} > \hat{p}_{T-1}]E[u(p_{T-1} - p_0)|p_{T-1} > \hat{p}_{T-1}] + \Pr[p_{T-1} < p_{liq}]u_{liq}$$
$$+ \Pr[p_{liq} \le p_{T-1} \le \hat{p}_{T-1}]E[u(v - p_0)|p_{liq} \le p_{T-1} \le \hat{p}_{T-1}]$$

where the indifference condition $u(\hat{p}_{T-1} - p_0) = E[u(v - p_0)]$ determines the cutoff $\hat{p}_{T-1}$. The liquidation price is p_{liq} and the expected utility in case of liquidation is u_{liq}. Prices posted by the MM depend on the history h_{T-2} (and the resulting inventory imbalance). If $p_{T-1} > \hat{p}_{T-1}$, then the arbitrageur closes his position. Going backwards, at the end of period T-3, the expected value of holding the position and possibly selling for a gain in $t=T$-2 is

$$D^{sell}_{T-2}(h_{T-3}) = \Pr[p_{T-2} > \hat{p}_{T-2}]E[u(p_{T-2} - p_0)|p_{T-2} > \hat{p}_{T-2}] + \Pr[p_{T-2} < p_{liq}]u_{liq}$$
$$+ \Pr[p_{liq} \le p_{T-2} \le \hat{p}_{T-1}]E[D^{sell}_{T-1}(h_{T-2})|p_{liq} \le p_{T-2} \le \hat{p}_{T-2}]$$

where the indifference condition $u(\hat{p}_{T-2} - p_0) = E[D^{sell}_{T-1}(h_{T-2})]$ determines the cutoff $\hat{p}_{T-2}$. Because $u(p_{T-2} - p_0) > u(0)$, the option value of any later period, in expectation, must be lower than the current period. Specifically, given the same price in period t and period $t+1$, the value of acquiring the position in $t+1$ is lower since there are fewer periods for the price to improve.

In period $t=0$, the arbitrageur determines whether $E[D^{sell}_1(h_0)] > u(0)$. If so, then the option to sell in future periods is valuable in expectation and he buys the asset at price p_0. An arbitrageur who owns the asset in period t determines whether to sell or hold. If $E[D^{sell}_{t+1}(h_t)] > u(p_t - p_0)$, he continues to hold, otherwise, he sells his position. The sale may be at a loss if $u(0) > u(p_t - p_0) \ge E[D^{sell}_{t+1}(h_t)]$ or at a profit otherwise. Note that this strategy implies that the arbitrageur may optimally continue to hold the asset even when he can realize an immediate cash profit by selling.

To see that trend trading is a profitable strategy over a generic space of parameters, consider the corner case of a risk-neutral arbitrageur, $p_0 = \bar{v}$, $\Pr[p < p_{liq}] = 0$, and ϕ such that $E[x_t | x_t > \phi] > \varepsilon > 0$ so that price moves in response to some order flow realization. The risk-neutral arbitrageur bears no cost buying the asset at this price and holding, so the worst case outcome for him is the same expected utility as not trading. On the other hand, he can sell for a gain in period 1 with positive probability. The risk-averse MM adjusts his price upward on positive inventory[5], so a loose lower bound on the expected profitability of this strategy is given by

$$\Pr\{x_0 > \phi\} E[p_1 - \bar{v} | x_0 > \phi] > \delta > 0$$

Therefore, there exists a positive measure set of prices above fair value for which this strategy is profitable to a risk-neutral arbitrageur. There exists some amount of leverage such that $\Pr[p < p_{liq}] > 0$ and the strategy is still profitable. Also, there exists a positive measure set of risk-averse utility functions "close" (in some relevant sense) to a risk-neutral utility such that buying the asset at a price close to fair value is strictly better than not trading.

Note that the expected profitability of this strategy depends on the size of the price deviation posted by the MM in response to an inventory imbalance. All else equal, a relatively more risk-averse MM posts larger price deviations in response to order flow and thus creates greater incentive for arbitrageurs to trend trade.

2.3.2 COUNTER-TREND TRADING

Arbitrageurs can sell after observing a price increase. This counter-trend trading allows the MM to offload inventory risk; the arbitrageur is compensated for this risk by a price premium above fair value. As before, the arbitrageur also gains a valuable option to close out the position if prices revert.

Consider a sale of a unit of the asset in period 0 at price p_0. Given public history h_{T-2} at the end of period $T-2$, the expected gain of holding and possibly buying back for a gain in $t=T-1$ is

$$D_{T-1}^{buy}(h_{T-2}) = \Pr[p_{T-1} < \hat{p}_{T-1}] E[u(p_0 - p_{T-1}) | p_{T-1} < \hat{p}_{T-1}] + \Pr[p_{T-1} > p_{liq}] u_{liq}$$
$$+ \Pr[\hat{p}_{T-1} \le p_{T-1} \le p_{liq}] E[u(p_0 - v) | \hat{p}_{T-1} \le p_{T-1} \le p_{liq}]$$

where the indifference condition $u(p_0 - \hat{p}_{T-1}) = E[u(p_0 - \bar{v})]$ determines the cutoff $\hat{p}_{T-1}$. If $p_{T-1} < \hat{p}_{T-1}$, then the arbitrageur closes his position. Expected value at $T-2$ is

$$D_{T-2}^{buy}(h_{T-3}) = \Pr[p_{T-2} < \hat{p}_{T-2}] E[u(p_0 - p_{T-2}) | p_{T-2} < \hat{p}_{T-2}] + \Pr[p_{T-2} > p_{liq}] u_{liq}$$
$$+ \Pr[\hat{p}_{T-2} \le p_{T-2} \le p_{liq}] E[D_{T-1}^{buy}(h_{T-2}) | \hat{p}_{T-2} \le p_{T-2} \le p_{liq}]$$

where the indifference condition $u(p_0 - \hat{p}_{T-2}) = E[D_{T-1}^{buy}(h_{T-2})]$ determines the cutoff $\hat{p}_{T-2}$.

In period $t=0$, the arbitrageur determines whether $D_1^{buy}(h_0) > u(0)$. If so, then the option to close out the asset position in future periods is valuable in expectation and he sells the asset at

---

price p_0. An arbitrageur who has a negative inventory of the asset in period t determines whether to buy or hold. If $D_{t+1}^{buy}(h_t) > u(p_0 - p_t)$, he continues to hold, otherwise, he closes his position. Note that the option value in this strategy implies that the arbitrageur may optimally continue to hold the asset even when price dips below fair value—in effect, the arbitrageur becomes a "trend-trader" for a time.

The portion of the parameter space where counter-trend trading is profitable is positive measure. Consider a risk-neutral arbitrageur, $\Pr[p > p_{liq}] = 0$, and price $p_0 > \bar{v}$. Clearly, the arbitrageur realizes an expected utility gain of $u(p_0 - \bar{v})$ just by selling the asset at this price and holding until time T. The option value of holding the position and closing it if price dips below $\bar{v}$ only adds to the value of this trade. As before, there exist "close" risk-averse utility functions where this strategy is profitable. Also, there exists some level of leverage such that $\Pr[p > p_{liq}] > 0$ and the strategy is profitable.

Whether counter-trend or momentum trading is a dominant action at a given period depends on the relative risk aversion of the MM. Consider that if risk aversion increases with inventory, price set by the MM may change in a convex manner. If the probability of an increase in MM inventory is close to the probability of a decrease in MM inventory (that is, ϕ is small), then it is possible to construct an example where the expected payoff to the arbitrageur is greater from momentum trading.

2.3.3 COLLATERAL CONSTRAINTS AND LEVERAGE

The MM has unlimited capital and does not face collateral constraints. Arbitrageurs have finite collateral and are therefore collaterally constrained and cannot generate order flow of arbitrary size. Assume also that arbitrageurs are sometimes leveraged. While this is an assumption, there are a number of justifications for this claim:

1. **Restrictions on fair value:** It is not optimal to provide collateral which is never drawn upon in any state of the world. Since the probability of a major currency becoming worthless in a short time period seems extremely close to zero, it may be optimal for arbitrageurs to be leveraged when trading currency assets if utility of total liquidation is not $-\infty$. Of course, the amount of leverage depends on the expected probability distribution of currency price.

2. **Profitable leveraged option:** Forced liquidations such as margin calls provide an arbitrageur with limited liability for losses, limited to the collateral in the account. For a martingale price process, the probability that the price does not reach the liquidation barrier before the information event is bounded away from zero even for infinite leverage (*i.e.*, no collateral). A price process as the one in the model allows positive expected payoff to a risk-neutral arbitrageur and he optimally chooses maximum leverage. Therefore, arbitrageurs with a not-too-high level of risk-aversion (and utility of total

liquidation greater than $-\infty$) optimally choose some amount of leverage which exposes them to liquidation risk.

3. **Lifecycle and borrowing constraints:** When viewed in a lifecycle model where the currency trading account is a small portion of lifetime wealth, the incentive to leverage the bet becomes stronger (as in Hervé, Tompaidis, and Yang (2013)). Because arbitrageurs cannot borrow against future earnings/wealth, it is straightforward to construct an example where it is optimal to make leveraged trades even for relatively high levels of risk-aversion. Since, in expectation, wealth may increase geometrically for a series of leveraged trades, a "young" arbitrageur has an incentive to leverage more than an "old" arbitrageur with comparable wealth (who may not have enough future income to offset possible losses).

4. **Biases:** Of course, in addition to rationally optimal reasons for leverage, behavioral biases may cause arbitrageurs to take on high leverage in excess of optimal limits.

In practice, retail FX trading allows high leverage on the order of 2.5% margin for major currencies for retail traders (and even higher for currency futures trading). Retail, as well as professional, traders are generally highly leveraged in their trades.

2.4 EQUILIBRIUM

Since counter-trend trading by arbitrageurs seems a more likely outcome in real-world situations, the following discussion restricts attention to those types of equilibria—that is, equilibria where arbitrageurs, in aggregate, only take positions toward fair value. Because of symmetry, it is sufficient to describe behavior in the top half of the price range.

The analysis herein treats arbitrageurs in aggregate. Individual arbitrageurs may engage in different strategies. In fact, some arbitrageurs may be trend-trading. Arbitrageurs may engage in inter-arbitrageur competitive behavior.

State variables in equilibrium are (a) N, the inventory held by the MM, (2) c, the cash (in the lead currency held by the MM, (3) A, the aggregate arbitrageur position size, and (4) K, the aggregate collateral (in the lead currency) of the arbitrageurs. Equilibrium consists of:
   a) $p_t(N, c, A, K)$, the MM's price-setting strategy,
   b) $b_t(N, c, A, K)$, the MM's inter-dealer trade strategy, and
   c) $a_t(p, N, c, A, K)$, the arbitrageurs' trade strategy
such that these strategies solve the arbitrageurs' and MM's problems.

The two main aspects of the examined equilibria are:
   1) the MM shifts price higher as arbitrageurs accumulate inventory and

2) the MM maintains a "shadow" inventory against arbitrageur inventory rather than trade it away into the inter-dealer market as he does for noise-trader order flow.

Higher prices incentivize arbitrageurs to take on larger positions, so a small bubble grows endogenously. Though the shocks driving the price process are martingale, some equilibrium outcomes exhibit price momentum.

The MM's problem in period *T-1*, where arbitrageurs react with order flow $a(p)$ is

$$\max_{p,b} \int_0^\infty f(v) \int g(x) U\big(v\{N_{T-1} - x - a(p) + b\} + c_{T-1} + xp + a(p)min(p, p_{liq})$$
$$- \bar{v}b\big) dxdv$$
$$s.t. \ |b| < \phi$$

where p_{liq} is the price at which the arbitrageur is totally liquidated. Due to limited liability, the MM cannot receive a price higher than p_{liq} in this margin call liquidation. At time *T-1* and price p_{T-1}, an arbitrageur who at time *T-2* had position A and collateral K (as measured in the lead currency) solves the problem

$$a(p; A, K, T-1) = \underset{a}{\operatorname{argmax}} E\big[u(v(A + \Delta A) + (v - p_{T-1})a + K + \Delta K)\big]$$
$$s.t. \Delta K = max\big((p_{T-1} - p_{T-2})A, -K\big),$$
$$\Delta A = I\big(|A|p_{T-1} > \gamma(K + \Delta K)\big)\left(|A| - \frac{\gamma(K + \Delta K)}{p_{T-1}}\right),$$
$$|A + \Delta A + a|p_{T-1} \leq \gamma[K + \Delta K]$$

Maximum allowed leverage is given by exogenous parameter γ. The term ΔA represents the forced liquidation (if any) due to margin constraints. The term ΔK represents the change in collateral due to price change subject to limited liability. For $p_{T-1} > \bar{v}$, a larger (negative) position size strictly increases expected return as price p_{T-1} increases. However, variance of return also strictly increases in position size strictly decreasing payoff to the risk-averse arbitrageur. Thus, $a(p; A, K, T-1)$ is well-defined.

Figure 1 illustrates the characteristics of response function $a(p; A, K, T-1)$ when the utility function has non-increasing risk-aversion[6]. For low prices (region A in the figure), an arbitrageur may benefit from "flipping" his position—closing out his existing position and taking a position in the opposite direction. Lower price allows the arbitrageur to take a larger position for the same collateral—a consequence of collateral being in the lead currency. Also, the arbitrageur's (short) position generates profit at lower prices and thus increases his collateral, allowing a larger position. Lower price also implies a greater deviation from fair value and higher expected return to a positive position. These effects combine to generate a decreasing $a(p)$ on region A. At fair value, the arbitrageur optimally takes a zero net position[7]. As price increases above fair value

---

[6] Increasing risk-aversion may have a local maximum for *a(p)* as price drops and the arbitrageur gains wealth.
[7] On a region around fair value, certain utility functions and distributions may make holding a position until the information event more costly than the benefit from taking a position on a slight price deviation. The arbitrageur

(region B), the arbitrageur finds it optimal to hold a short position. The greater the price deviation, the greater the expected benefit, so the arbitrageur increases his desired position size, net of his initial position A. Because the arbitrageur holds a short position at the start of the period, higher prices reduce his collateral and may reduce optimal position size. The response function $a(p)$ has a minimum due to limited collateral. When prices get sufficiently high (region C), the arbitrageur may reduce his position size due to risk aversion or may have his position reduced via forced liquidations from margin calls. Note that even a risk-neutral arbitrageur who takes on maximum allowed leverage has position $|A| = \gamma K/p$ which decreases in price (ignoring the additional reduction in collateral from his existing losing position). This consequence, again, results from collateral being in the lead currency.
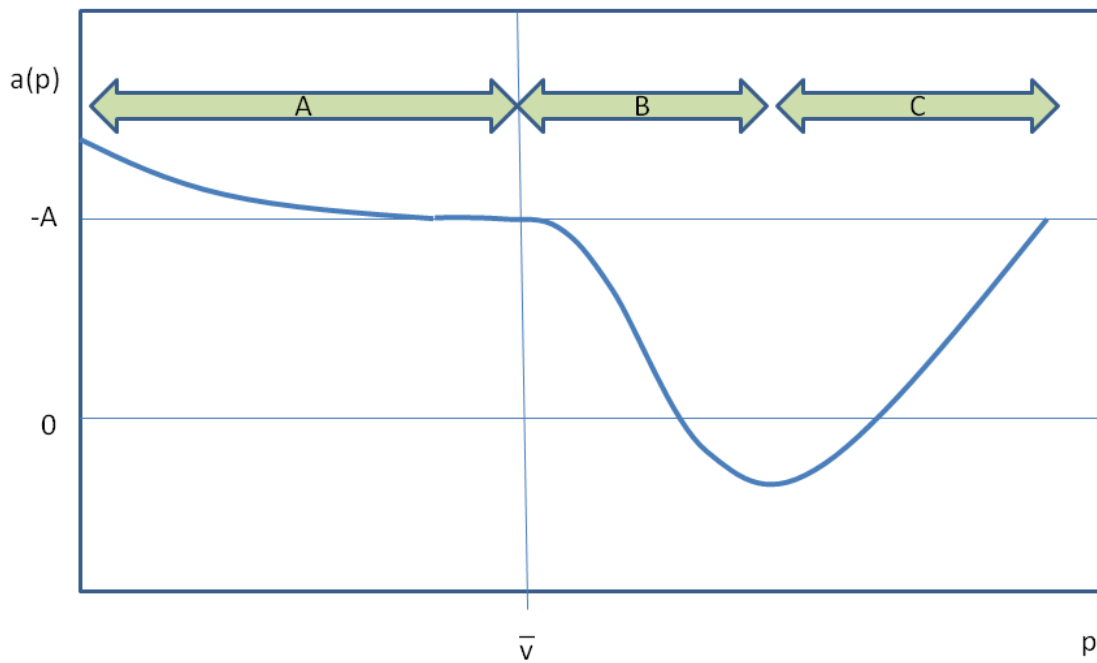

Figure 1: Arbitrageur response function a(p)

Consider an earlier period when arbitrageurs have no (or little) inventory. If the MM sets price sufficiently above fair value, arbitrageurs optimally sell the asset. The MM anticipates this order flow and may benefit on the costs of inventory control by reducing the posted price below the level of the no-arbitrageur case. Looking to period *T-1* after arbitrageurs have accumulated a negative position, the MM expects order flow *a(p)*. His price setting response depends on his

---

would then prefer a zero net position over that price region. In other words, there may exist $q > \bar{v}$ such that $E[u(q - \bar{v})] < u(0)$.

expected inventory at the end of the period, but also on any surplus he can extract from the arbitrageurs.

2.4.1 PRICE

When arbitrageurs are a very small part of the market, the MM does not gain from adjusting price too much from the no-arbitrageur case. As the relative size of arbitrageur order flow increases, the MM's pricing power also increases. When arbitrageurs hold a position, the MM has market power in his pricing depending on the relative size of the position. Consider that if there were no noise traders, the MM can set price to extract full surplus from the arbitrageur where the arbitrageur's reservation utility is to hold the asset until the information event. If the arbitrageur is leveraged, the MM may also set price to liquidate the arbitrageur and thereby extract not just surplus but the arbitrageur's wealth. In a subgame perfect equilibrium, the MM cannot commit not to exploit the arbitrageur, and the market breaks down. In order for arbitrageurs to enter the market and for there to exist a non-trivial subgame perfect equilibrium, the price set by the MM must, in expectation, offer non-negative surplus to the arbitrageurs. A necessary condition for such an equilibrium depends on the relative size of arbitrageur order flow versus noise trader order flow[8].

Multiple equilibria exist when arbitrageurs have positive expected surplus because expected surplus may be distributed between the two arbitrageur actions of selling and buying back. For instance, the MM may not adjust price much from the no-arbitrageur level when arbitrageurs are expected to take their initial position, but may set price higher to extract all remaining expected surplus (constrained by the MM's expectation of noise trader order flow) when arbitrageurs buy back. This particular strategy is subgame perfect since the MM cannot commit to any other price setting strategy.

In a subgame perfect equilibrium, one may expect price deviation (from the no-arbitrageur case) to increase in the size of arbitrageurs' positions. To gain intuition on this price setting behavior, set $\phi = 0$ to ignore inter-dealer trade. Total liquidation of the arbitrageur is a corner case, which can be handled separately, so the price received from the arbitrageur is taken to be p—that is, no liquidation barrier. In period T-1, the MM's problem is

$$\max_{p} \int_{0}^{\infty} f(v) \int g(x) U\big(v\big(N_{T-1} - x - a(p)\big) + xp + a(p)p\big) \, dx \, dv$$

where cash is taken to be zero. The utility and distribution functions are continuous and $a(p)$ is piece-wise continuous, so it is possible to take derivatives through the integral. First-order conditions give:

$$E\big[U'\big(v\big(N_{T-1} - x - a(p)\big) + xp + a(p)p\big) \cdot \big(x - va'(p) + a(p) + pa'(p)\big)\big] = 0$$

---

[8] It is likely possible to construct repeated games with reputation where a trade equilibrium exists and does not exist in the one-shot game described here.

If the MM is risk-neutral, the solution is implicitly given by $(p - \bar{v})a'(p) + a(p) = 0$. Since the MM is risk-averse, the optimal price is lower (to reduce variance). In the risk-neutral scenario, it is clear that if $a'(p) > 0$, then there is no local maximum and only a corner solution where p is the total liquidation price. This result is the max for the MM's problem under the assumption of risk-neutrality and demonstrates the incentive to the MM to force liquidation when arbitrageurs hold any inventory. Risk-aversion restrains the MM to a range of prices. If risk-aversion effectively holds the MM to the range of price where $a'(p) < 0$ and $a(p) > 0$, then the solutions is that single local maximum. The FOC solution point where $a'(p) > 0$ and $a(p) < 0$ is a minimum. In particular, if the MM has zero inventory and if arbitrageurs sell on any price deviation, that is, $a(p) < 0$, then the MM has no incentive to raise price above $\bar{v}$.

Since the option value of the arbitrageur's position decreases in time, the incentive to hold or increase a position also decreases at every price. Therefore, $a(p)$ decreases in t over region A and increases over regions B and C. The point/region where $a(p) = -A$ shifts right since earlier periods have more option value, so the arbitrageur may be willing to hold a non-zero position even when the posted price is the fair value price. Smaller collateral—that is, an increase in leverage as position size remains the same—increases the probability of liquidation while the expected payoff remains the same. Thus, the incentive to hold or increase a position decreases at every price. The arbitrageur response $a(p)$ decreases in K over region A and increases over regions B and C. Since position liquidation occurs at a faster rate with higher leverage, the slope of $a(p)$ increases over region C. Increasing initial position size has a similar effect as decreasing collateral since leverage increases. However, $a(p)$ goes up overall as region A must increase also. This change occurs since there is more inventory—for instance, $a(\bar{v}; A, K, T - 1) = -A$.

The MM's response to higher $a(p)$ is to post higher price. As time goes by and arbitrageurs accumulate a larger position, the MM's posted prices move higher in expectation. This effect initially feeds on itself. Higher price causes arbitrageurs who do not have too much leverage to take on a short position. As in Kondor (2009), arbitrageurs optimally expect to allocate collateral across multiple periods, so their aggregate position size grows gradually. The resulting price-history displays momentum.

Since the MM adjusts prices only slightly at low levels of arbitrageur inventory, price dynamics closely resemble a martingale process. A more substantial price deviation may cause arbitrageurs to take on larger positions which in turn cause the MM to shift price higher. This process drives the inflation of a small price bubble.

It is important to reiterate that the described outcomes deal with arbitrageurs in aggregate. Individual arbitrageurs may, in fact, trade with a momentum strategy. The MM reacts based on the overall level of arbitrageur inventory. Individual traders who attempt to trade with

momentum reduce the inventory aggregate and the growth of the bubble. This risk prevents arbitrageurs from consistently profiting from these dynamics.

2.4.2 INVENTORY

The MM decides how much inventory to offload into the inter-dealer market. Since any inventory represents a risk and expected utility loss to the MM, one may suppose that the MM maximizes his inter-dealer trade as in the no-arbitrageur case. Certainly, in period *T-1*, the MM sets *b* to reduce his inventory as close to zero as possible. However, in earlier periods, the overhang of arbitrageur inventory may hit the market at any time. The MM can do strictly better by keeping opposite sign inventory to absorb this expected order flow.

Suppose that arbitrageurs acquire (negative) inventory from the MM. The MM's problem in inventory management may be characterized in the following manner. There are three possible outcomes:

(1) The price path is such that a price below fair value is realized and arbitrageurs close their positions. This price occurs earlier than a price above fair value where arbitrageurs close their positions.
(2) The price path is such that price remains above the region where arbitrageurs close their position and arbitrageurs wind up holding inventory until the information event.
(3) The price path is such that a price above fair value is realized and arbitrageurs close their position. This price occurs earlier than a price below fair value where arbitrageurs close their positions.

The MM can trade the inventory acquired from arbitrageur order flow to the inter-dealer market at fair value or he can hold it. If outcome (2) is realized, the MM would have been strictly better off trading the inventory to the inter-dealer market because the certainty equivalent fair price received there is better for a risk-averse agent than holding risky inventory. If outcome (1) is realized, the MM takes a cash loss on his inventory because he could have traded the inventory at fair value. If he trades inventory to the inter-dealer market, the MM absorbs this arbitrageur order flow (from the buy back trade) at a lower cost-basis—he effectively skips the loss between fair value and this other lower price. Only in outcome (3) does the MM do better by holding inventory. He may take a cash loss if arbitrageurs close the position at a lower price than the price at which they opened the position, but the MM's loss is smaller than the alternative loss of trading at fair value. It may also happen that arbitrageurs close the position above the original price and thus generate a cash gain for the MM.

Therefore, the MM considers the probability and expected losses from outcomes (1) and (2) versus the probability and expected gains from outcome (3). Two factors affect the MM's decision: time until the information event and the amount of leverage used by the arbitrageurs. For arbitrageurs, the option value of holding a position decreases as time approaches the information event. Therefore, arbitrageurs close their positions at higher prices. Leverage lowers

the price at which an arbitrageur closes his position. Moreover, when arbitrageurs are leveraged and price becomes sufficiently high, collateral constraints cause arbitrageurs to close positions for a cash loss. These characteristics of the arbitrageur response $a(p)$ were described in section 2.4.1.

Suppose the MM has total inventory of zero. The MM's problem for period *T-2* in inventory control may be written as:

$$\max_b \int_0^\infty f(v) \int g(x) U \left(v\{-x - a_{T-1}(p_{T-1}^*) + b_{T-1}^* + b\} + c_{T-1} + xp_{T-1}^* \right.$$
$$\left. + a_{T-1}(p_{T-1}^*)min(p_{T-1}^*, p_{liq}) - \bar{v}(b + b_{T-1}^*) \right) dxdv$$
$$s.t. \ |b| < \phi$$

where p_{T-1}^* is the optimal price chosen for period *T-1*, a_{T-1} is the arbitrageur response in period *T-1*, and b_{T-1}^* is the optimal inter-dealer trade size in period *T-1*. Price p_{T-1}^* is set by the MM based on inventory at the start of *T-1*. However, in equilibrium, the MM knows that inventory will be changed by $a_{T-1}(p_{T-1}^*)$ regardless of noise trader order flow. This quantity is non-stochastic at the end of period *T-2*. Setting $b = a_{T-1}(p_{T-1}^*)$ puts the MM's (stochastic) inventory at $N_{T-1} = -x + b_{T-1}^*$, which has expected value of zero and is the amount of inventory which minimizes utility loss to the MM at the information event. In the described setup, the MM traded in the inter-dealer market to acquire a "shadow" to arbitrageur inventory. Optimally, the MM should have kept arbitrageur inventory at the time arbitrageurs opened their positions. If the MM had started the period with inventory $-A$, then the optimal choice of $b = A + a_{T-1}(p_{T-1}^*)$ represents trading into the inter-dealer market the amount of inventory arbitrageurs expect to hold into the information event.

For earlier periods, the MM optimally reduces shadow inventory as the probability increases of him having to post a price below fair value. If noise trader shocks are not to big, then the probability of the MM having to post a big price change are small. The MM is better off hedging the potential loss of outcome (1) by preemptively reducing arbitrageur shadow inventory. Note, however, that arbitrageurs voluntarily reduce their positions as price approaches fair value. This difference at close to fair value between optimal "shadow" inventory and arbitrageur positions becomes smaller in time since arbitrageurs reduce positions as option value decreases.

If arbitrageurs have some chance of holding inventory into the information event, the MM can always trade to the inter-dealer market in period *T-1*. Of course, if the expected amount of inventory arbitrageurs hold into the information event is large, then the total of arbitrageur inventory and noise trader order flow may exceed ϕ. In that case, the MM optimally trades some of his shadow inventory in earlier periods. Earlier periods require less of this hedging since there is greater chance of the MM's inventory being closer to zero by the time period *T-1* comes.

When *T* is far away, arbitrageurs expect to profitably close their positions prior to the information event. If noise trader order flows are not too extreme, price ought to change somewhat smoothly. The prices at which arbitrageurs profitably close positions are relatively close to fair value. Thus, a reasonable heuristic rule for the above analysis may be for the MM to keep all arbitrageur order flow as shadow inventory, especially when arbitrageurs become more leveraged.

2.4.3 REMARKS

A simplifying assumption was that aggregate arbitrageur inventory is public information. This assumption is not true in the real-world. If arbitrageur positions are not public information, an individual arbitrageur is at a bigger disadvantage. He must estimate market leverage based on other public information—a complex task. The MM, who does have this information, then extracts more surplus from arbitrageur trades by shifting price higher than in the case of public information. Arbitrageurs do not know exactly how much of any price change to ascribe to noise trader order flow and how much to surplus extraction. Their response is based on a probabilistic estimate of MM actions. Asymmetric information may then reduce trade.

Traders who have access to non-public information on arbitrageur positions have an advantage in this market. Such traders may see the formation of price momentum due to collateral strain and trade in that direction. As arbitrageurs begin to close their positions, the bubble begins to deflate and the informed trader exits his trade. As a heuristic, an informed trader may simply follow a representative arbitrageur. The informed trader waits until the arbitrageur becomes leveraged and buys when the arbitrageur sells. When the uninformed arbitrageur buys back his position or is liquidated, the informed trader sells his position.

The described mispricing cannot be arbitraged away so long as the amount of collateral in the market cannot change quickly. Suppose an arbitrageur has two accounts—each with a different market maker. This arbitrageur can sell the asset in one market and buy it in the other if the first market has a high price. However, this strategy does not qualitatively affect the described equilibrium because the arbitrageur cannot transfer collateral from one market to the other. Each individual market behaves as before.

3 EVIDENCE FROM PRICE DATA

According to the model, small inventory imbalances from noise trader order flow and the resulting small price deviations ought to be corrected fairly quickly by the MM's trade in the inter-dealer market. On the other hand, a larger imbalance in a market with collaterally strained—that is, highly leveraged—market participants should result in price momentum. A collaterally strained market which has recently seen price momentum exhibits greater fragility in being able to absorb random order flow.

The FX market is decentralized; no central exchange tracks trades, though some inter-dealer trade data are made publically available. It is a 24 hour a day market. The examined dataset consists of all posted bids and offers from one particular dealer for the time period March 15, 2012 to January 24, 2013. This dataset allows the examination of price-setting by the dealer—transaction data would not give a complete picture. Analysis of the model is substantially limited, however, by the lack of data on traders and their level of leverage and the lack of data on information events. Using available data, evidence to suggest for the model comes from price movements for momentum and reversion effects at time-scales which correspond to dynamics in the model. In particular, price ought to revert for small price moves due to MM inter-dealer trade and trade from arbitrageurs when collateral is not strained. After a larger price move that does not correspond to an information event, the model suggests that arbitrageurs counter-trade and deplete their collateral causing price to continue to rise. At a longer timescale, this momentum effect ought to revert.

Empirical analysis here focuses on the EUR-USD currency pair (euro vs. US dollar) as it is the most actively traded FX instrument. This price series in the dataset has over 8.5 million data points. With prices measured by the midpoint of the spread, the euro begins the period at 1.3104 dollars and ends at 1.3368—a euro appreciation of 2.01%. The mean is approximately 1.2833, the maximum is 1.3404, and the minimum is 1.2042. Other currency pairs are examined later.

3.1 Autoregressive conditional heteroskedastic (ARCH) model

When looking at financial time series with possible momentum/reversion effects, researchers often use an autoregression model where the variance of error terms evolves from history. This type of econometric model provides a useful baseline.

Measuring price momentum (or reversion) via auto-regressive models involves some guesswork as to the correct timescale to use. Moreover, it may be the case that the timescale is not constant. For instance, it may be that momentum exists at lower frequencies during normal times, but does not manifest at those frequencies during stressed periods (when it may show up at higher frequencies). ARCH models attempt to account for this effect by modeling changes in variance through time. However, the question of timescale persists.

A simple ARCH(1) model looking at twenty lagged auto-regressive terms for a difference log series on the price of EUR-USD sampled every 30 minutes finds some evidence of auto-correlation. Coefficients are listed in Table 1 (lags 13 to 20 have p-values greater than 0.05 and are omitted):

Lag	Coefficient	p-value
1	-0.069570	0.0000
2	-0.003624	0.5847
3	-0.017147	0.0116
4	-4.31E-05	0.9952
5	0.013880	0.0545
6	-0.002900	0.7044
7	-0.017812	0.0126
8	-0.009318	0.1745
9	7.98E-05	0.9916
10	-0.009215	0.2162
11	-0.009619	0.1961
12	0.020128	0.0208

Table 1: Coefficients for ARCH(1) model with 20 lags (last seven omitted) for EUR-USD price series sampled at 30 minute frequency. (N=10578)

The largest effect appears to be reversion from a 30 minute lag. There may be a small momentum effect at 6 hours. Other effects (if significant) seem small. In terms of the proposed model, it may be that reversion is due to either arbitrageur counter-trend trading or due to the MM offloading a spike in inventory into the inter-dealer market. This regression model, however, does not examine any non-linear price response effects predicted by the presented market model.

3.2 FINDING NON-LINEAR PRICE REACTIONS

Looking directly at the conditional price distribution of EUR-USD over various timescales, one can see that momentum effects are more pronounced at higher frequencies. Table 2 shows statistics for $r_{+\tau}$ given $r_{-\tau}$ where $r_{+\tau}$ is the percent price move between the current sample time to a time τ minutes ahead and r_{-T} is the price move between the current sample time and a time τ minutes behind. The price is the midpoint of the posted bid-ask spread sampled at 10 minute intervals. Tests[9] for normality fail for all examined conditional distributions. Price data in the series are in increments of 10^{-5} euros, and the analysis rounds percent moves to a hundredth of a basis point (that is, in 0.0001% increments). In Table 2, ranges for r_{-T} are chosen to divide the sample into "small" moves and "big" moves with "small" moves defined to be roughly 2/3 of the sample. A linear regression of price responses on past moves provides an estimate of momentum or reversion responses.

τ	$r_{-\tau}$	Summary statistics			Coefficient (std. error) [p-value]
		Mean (StdDev) Range and sample size	Median, Skewness	Kurtosis	Constant, Lag price move
10 mins.	Unconditioned	0.7386 (435.8107) Range=7490 to -4212, N=31200	0, 0.2728	15.9463	] 0.734175 (2.467293) [0.7660] 0.005866 (0.005662) [0.3001]
	$-300 \leq r_{-\tau} \leq 300$ (-0.03% to 0.03%)	1.315374 (373.5521) Range=5062 to -4212, N=20899	0, -0.0519	15.3215	1.315362 (2.583471) [0.6107] -0.051979 (0.017176) [0.0025]
	$r_{-\tau} < -300$ (less than -0.03%)	4.8829 (532.4003) Range=4000 to -4141, N=5157	22, -0.5219	10.0784	28.83758 (14.13325) [0.0414] 0.039133 (0.019659) [0.0466]
	$r_{-\tau} > 300$ (more than 0.03%)	-5.7595 (548.5817) Range=7490 to -3142, N=5144	-40, 1.3829	16.2654	-54.23974 (14.06961) [0.0001] 0.078416 (0.019113) [0.0000]

---

[9] Econometric normality tests available in eViews software (such as Lillefors, Cramer-von Mises, Watson, and Anderson-Darling) reject the hypothesis of a normal distribution.

30 mins.	Unconditioned	1.9891 (752.4765) Range=12247 to -11811, N=31196	0, 0.1711	16.3473	1.989731 (4.260408) [0.6405] -0.000441 (0.005642) [0.9377]
	$-500 \leq r_{-\tau} \leq 500$ (-0.05% to 0.05%)	0.3022 (660.9965) Range=12247 to -6466 N=20660	0, 0.2677	17.7034	0.392802 (4.597877) [0.9319] -0.055272 (0.018254) [0.0025]
	$r_{-\tau} < -500$ (less than -0.05%)	22.2355 (891.6154) Range=5896 to -6659, N=5287	70, -0.4488	8.5823	93.69116 (22.91083) [0.0000] 0.068314 (0.018511) [0.0002]
	$r_{-\tau} > 500$ (more than 0.05%)	-11.7643 (918.9248) Range=10859 to -11811, N=5249	-77, 0.5727	17.4453	-51.24506 (23.06745) [0.0264] 0.037401 (0.018255) [0.0405]
4 hours	Unconditioned	18.9953 (2069.061) Range=15679 to -15410, N=31154	0, 0.2161	7.1796	19.10061 (11.72134) [0.1032] -0.014807 (0.005615) [0.0084]
	$-1400 \leq r_{-\tau} \leq 1400$ (-0.14% to 0.14%)	-11.5550 (1993.761) Range=14765 to -11369, N=19623	0, 0.2095	7.3542	-12.13338 (14.23921) [0.3942] 0.026310 (0.019842) [0.1849]
	$r_{-\tau} < -1400$ (less than -0.14%)	146.6796 (2120.576) Range=15679 to -10985, N=5871	209, 0.2684	6.8688	86.02038 (61.02208) [0.1587] -0.021158 (0.018970) [0.2647]
	$r_{-\tau} > 1400$ (more than 0.14%)	-7.5327 (2258.003) Range=11582 to -15410, N=5882	-117, 0.1742	6.7860	-115.7071 (60.28569) [0.0550] 0.036835 (0.017805) [0.0386]
24 hours	Unconditioned	-16.7209 (4853.885) Range=21154 to -17220, N=30916	-183, 0.3041	4.2190	-24.42802 (27.59337) [0.3760] 0.044555 (0.005508) [0.0000]
	$-4400 \leq r_{-\tau} \leq 4400$ (-0.44% to 0.44%)	-19.20150 (4854.621) Range=21154 to -17220, N=20164	-401, 0.4131	4.4335	-13.94148 (34.21190) [0.6836] 0.052361 (0.015120) [0.0005]
	$r_{-\tau} < -4400$ (less than -0.44%)	-274.6699 (5088.355) Range=19998 to -13292, N=4905	61, 0.2654	4.2608	1235.509 (229.4039) [0.0000] 0.205432 (0.029616) [0.0000]
	$r_{-\tau} > 4400$ (more than 0.44%)	208.2252 (4634.828) Range=15793 to -13693, N=4932	165, -0.0556	3.2894	-331.3975 (157.6611) [0.0356] 0.072648 (0.019598) [0.0002]

Table 2: Summary statistics for estimated conditional price distribution at various timescales. Prices are from EUR-USD series sampled every 10 minutes. Returns are represented rounded to a hundredth of a basis point (0.0001% increments). Last column shows regression results for forward price move (dependent variable) on constant and past price move.

The observed decrease in kurtosis as frequency decreases may come from the underlying price-generating process: a noise process with occasional large jumps. Large jumps overshadow the noise diffusion process at higher frequencies. It also may be that the MM can absorb small amounts of inventory without adjusting prices too much, but a sustained large increase has a non-linear effect on his posted prices.

The response of prices to previous moves appears to have timescale dependency. At the 24 hour frequency, small moves have momentum, big moves in dollar strength have some reversion, and big moves in euro strength have little clear direction. Both types of big moves appear to have momentum as they get even bigger. Very big moves in dollar strength, especially, show high momentum[10] as the coefficient is 0.2054. At 4 hours, the picture becomes more muddled as there may be some weak momentum for small moves, weak reversion for big moves when the dollar strengthens, and reversion for big moves when the euro strengthens. Even bigger moves, especially for euro strength, appear to have momentum. The switchover to momentum occurs somewhere above the chosen cutoff as the majority of price responses are mean reverting. At 30 minutes, more clear relationships emerge as small price moves appear to show reversion while

---

[10] There may be risk effects on various currencies. For instance, during the past several years, it has been generally considered that buying the euro is a "risk-on" trade—that is, a risk loving trade. One expects that big "risk-off" trades show momentum due to risk-aversion. The Japanese yen shows the reverse effect vs. the dollar as the yen has been a "risk-off" trade. Specifically, in a regression as above, the coefficient on USD-JPY responses to price moves below -0.44% is 0.0270 (p-value 0.3750) while the coefficient on responses to price moves above 0.44% is 0.1501 (p-value 0.0000).

big moves have momentum. At 10 minute frequency, reversion continues for small moves. Big moves have momentum with stronger responses when the euro strengthens.

According to the model, at the timescale at which the MM does not easily resolve inventory imbalances, prices should exhibit reversion for small moves and momentum for big moves. Therefore, examining the timescale on the order of a 10 to 30 minute frequency is appropriate from the data above and is likely consistent with the timescale of actors in the model. The top half of Table 3 shows regression results for price responses for more finely sliced regions of the 20 minute timescale. In line with model predictions, it seems that prices show mild reversion after price moves less than about 0.1% but then show momentum for larger moves, especially for moves greater than 0.3%.

On a longer timescale, a likely[11] reason for a reversion from a price move is that the MM resolves his inventory buildup. A timescale of about 4 hours (half a business day) seems to be a reasonable estimate to see reversion from inventory balancing. Since the euro appreciated by about 201 basis points during the sample period, there is an average increase of about 0.6381 basis points per day, so four hour price responses are de-trended by -0.11bp. The bottom half of Table 3 shows regression results for price responses for the 4 hour timescale. While small moves have definite reversion (and stronger than at the 20 minute response), larger moves have little clear direction.

Time: Lag, Response	$r_{-\tau}$ range Sample size	Coefficient (std. error) [p-value] Constant, Lag price move	Price response description and Est. size of response (bp. EUR)
	$r_{-\tau}$ <-3000 (less than -0.30%) N=56	940.5538 (863.0737) [0.2807] 0.338360 (0.218293) [0.1270]	Likely momentum >0.75bp
	-3000≤ $r_{-\tau}$<-1500 (-0.30% to -0.15%) N=398	-88.64667 (272.0044) [0.7447] 0.003761 (0.138362) [0.9783]	Uncertain Possible momentum 0.82bp to 0.77bp
	-1500≤ $r_{-\tau}$<-500 (-0.15% to -0.05%) N=3767	84.00615 (39.64199) [0.0341] 0.071507 (0.046904) [0.1275]	Reversion 0.48bp to Likely momentum 0.23bp
20 min., 20 min.	| $r_{-\tau}$ |≤500 (less than 0.05%) N=22814	-0.257640 (3.548835) [0.9421] -0.056539 (0.014674) [0.0001]	Reversion < 0.28bp
	500< $r_{-\tau}$ ≤1500 (0.05% to 0.15%) N=3663	-77.91230 (40.40698) [0.0539] 0.070383 (0.047795) [0.1409]	Reversion 0.43bp to Likely momentum 0.28bp
	1500< $r_{-\tau}$ ≤3000 (0.15% to 0.30%) N=450	53.99474 (266.8396) [0.8397] 0.029692 (0.134710) [0.8256]	Uncertain Possible momentum 0.99bp to 1.43bp
	$r_{-\tau}$ >3000 (more than 0.30%) N=50	1035.424 (751.4418) [0.1746] -0.143954 (0.170644) [0.4031]	Likely momentum <6.03bp
	$r_{-\tau}$ <-3000 (less than -0.30%) N=56	3419.027 (1535.429) [0.0302] 0.941515 (0.388349) [0.0187]	Momentum >6.06bp
20 min., 4 hours	-3000≤ $r_{-\tau}$<-1500 (-0.30% to -0.15%) N=398	-1279.576 (731.4460) [0.0810] -0.722244 (0.372068) [0.0529]	Likely momentum 2.07bp to Likely reversion 8.76bp
	-1500≤ $r_{-\tau}$ <-500 (-0.15% to -0.05%) N=3763	88.05525 (128.5713) [0.4935] 0.052741 (0.152171) [0.7289]	Uncertain Possible reversion <0.50bp
	| $r_{-\tau}$ |≤500 (less than 0.05%)	-5.431845 (12.91387) [0.6740]	Reversion < 0.88bp

---

[11] In the described model, one expects to see a reversion effect after a large momentum move which is followed by an information event. A price distortion away from fair value skews price moves at the information event so that half of them are larger than without the distortion. Since the quadratic loss function in the regression puts greater weight on large moves, larger distortions imply a greater measured reversion. However, it seems unlikely that information events occur sufficiently often for this type of effect to materially impact results.

	Coefficient (std. error) [p-value]	Price response description
N=22807	-0.143326 (0.053397) [0.0073]	
500< $r_{-\tau}$ ≤1500 (0.05% to 0.15%) N=3652	4.124921 (130.6237) [0.9748] 0.026492 (0.154630) [0.8640]	Uncertain Possible momentum 0.06bp to 0.32bp
1500< $r_{-\tau}$ ≤3000 (0.15% to 0.30%) N=450	327.6791 (684.7741) [0.6325] -0.089673 (0.345699) [0.7954]	Uncertain Possible momentum 1.82bp to 0.48bp
$r_{-\tau}$ >3000 (more than 0.30%) N=50	1333.331 (1988.939) [0.5058] -0.167147 (0.451665) [0.7130]	Uncertain Possible momentum <8.21bp

Table 3: Regression results for price response from lagged price move for EUR-USD currency pair separated by size of price move for a timescale where the market maker likely does not resolve inventory shocks (20 mins.) and a time scale where he has likely been able to resolve his inventory imbalance (4 hours). Response sizes (in right column) for 4 hours are de-trended by -0.11 bp.

Measuring reversion after momentum at the 4 hour timescale poses two issues. One is that since momentum begets more momentum it is likely that price does not revert for a while. Supposing that price does revert at some point raises the issue of measurement since the response period includes both a trend and a trend reversal and one should see no clear direction in the price response. Looking at a longer lagged price move may reduce these two effects. In Table 4, a four hours lagged price move tends to revert except for large moves of euro strength of 0.3-0.4% or more.

Time: Lag, Response	$r_{-\tau}$ range Sample size	Coefficient (std. error) [p-value] Constant, Lag price move	Price response description and Est. size of response (bp. EUR)
	$r_{-\tau}$ <-4000 (less than -0.4%) N=990	36.79082 (253.7199) [0.8847] -0.015111 (0.044904) [0.7365]	Uncertain Possible reversion >0.88bp
	-4000≤ $r_{-\tau}$ <-2000 (-0.4% to -0.2%) N=2953	285.4087 (195.3130) [0.1440] 0.015175 (0.069027) [0.8260]	Likely reversion 2.44bp to 2.14bp
	-2000≤ $r_{-\tau}$ <-1000 (-0.2% to -0.1%) N=3824	-415.8666 (181.4730) [0.0220] -0.268396 (0.123714) [0.0301]	Momentum 1.59bp to Reversion 1.10bp
4 hours, 4 hours	| $r_{-\tau}$ |≤1000 (less than 0.10%) N=15694	-4.708660 (15.79723) [0.7657] 0.030810 (0.029533) [0.2969]	Uncertain Possible momentum < 0.46bp
	1000< $r_{-\tau}$ ≤2000 (0.1% to 0.2%) N=4048	183.5300 (162.8077) [0.2597] -0.157919 (0.111461) [0.1566]	Likely momentum 0.15bp to Likely reversion 1.44bp
	2000< $r_{-\tau}$ ≤4000 (0.2% to 0.4%) N=2618	-924.5154 (235.6412) [0.0001] 0.319837 (0.083231) [0.0001]	Reversion 2.96bp to Momentum 3.44bp
	$r_{-\tau}$ >4000 (more than 0.4%) N=1022	979.2915 (238.5988) [0.0000] -0.133654 (0.039051) [0.0006]	Momentum < 4.33bp

Table 4: Regression results for price response from lagged price move for EUR-USD currency pair separated by size of price move for a timescale where short-term momentum is likely attentuated (4 hours). Response sizes (in right column) are de-trended by -0.11 bp.

The time required for an MM to resolve inventory shocks may depend on factors such as the volatility of prices, the size of that market, macro trends and economic environment, and elasticity of supply and demand in the inter-dealer market. Therefore, the predicted effects may show up at different timescales for different currency pairs. However, one would expect major currencies to have somewhat comparable timescales.

Looking at the USD-CHF pair (US dollar vs. Swiss franc), the distribution of price moves for USD-CHF at a 20 minute frequency has a kurtosis of 27.0224 versus 17.2497 for EUR-USD. A more leptokurtic distribution means that there are relatively more small price moves for USD-CHF which may reduce inventory risk to a MM on average thus implying smaller responses and/or quicker reversion. Table 5 shows comparable measures and results for USD-CHF. The first price in the series is 0.92135 and the last one is 0.92938 as the dollar gained 0.872% for an average gain of 0.2768 basis points per day. The highest price is 0.99723 and the lowest is 0.90021. Table 5 shows the predicted reversion and momentum effects for the 20 minute timescale. At 4 hours, reversion is seen for small moves. Big moves may have reversion, though franc strength moves seem to have momentum. Looking at longer lagged price moves (to attenuate short-term momentum), it appears that reversion exists for larger moves except for moves of franc strength greater than 0.3%.

Time: Lag, Response	$\lvert r_{-\tau} \rvert$, Sample size	Coefficient (std. error) [p-value] Constant, Lag price move	Price response description Est. size of response (bp. USD)
20 min., 20 min.	$r_{-\tau}$ <-1000 (less than -0.10%) N=1290	13.58658 (69.46458) [0.8450] 0.019036 (0.040061) [0.6347]	Uncertain Possible momentum > 0.05bp
	-1000≤ $r_{-\tau}$<-500 (-0.10% to -0.05%) N=2801	153.6744 (66.50089) [0.0209] 0.167627 (0.093751) [0.0739]	Reversion 0.70bp to likely momentum 0.14bp
	$\lvert r_{-\tau} \rvert$<500 (less than 0.05%) N=21754	-1.085565 (3.680894) [0.7681] -0.079599 (0.015027) [0.0000]	Reversion 0.41bp
	500≤ $r_{-\tau}$ <1000 (0.05% to 0.10%) N=2810	125.6315 (68.50984) [0.0668] -0.233472 (0.096251) [0.0153]	Reversion <1.08bp
	$r_{-\tau}$ >1000 (more than 0.10%) N=1281	-105.9313 (60.82680) [0.0818] 0.096495 (0.035739) [0.0070]	Momentum > 0 bp
20 min., 4 hours	$r_{-\tau}$ <-1000 (less than -0.10%) N=1287	-232.1224 (183.9207) [0.2071] -0.051920 (0.106020) 0.6244]	Uncertain Possible momentum <1.85bp
	-1000≤ $r_{-\tau}$<-500 (-0.10% to -0.05%) N=2797	294.2240 (219.1580) [0.1795] 0.369681 (0.308981) [0.2316]	Uncertain Possible reversion 1.04bp to momentum 0.80bp
	$\lvert r_{-\tau} \rvert$<500 (less than 0.05%) N=21743	32.39733 (13.06379) [0.0131] -0.137448 (0.053333) [0.0100]	Reversion < 0.32bp to <1.05bp
	500≤ $r_{-\tau}$ <1000 (0.05% to 0.10%) N=2820	196.0975 (218.2868) [0.3691] -0.406990 (0.306960) [0.1850]	Likely reversion 0.13bp to 2.16bp
	$r_{-\tau}$ >1000 (more than 0.10%) N=1278	-193.4617 (168.9244) [0.2523] 0.054400 (0.099221) [0.5836]	Uncertain Possible reversion <1.44bp
4 hours, 4 hours	$r_{-\tau}$ <-3000 (less than -0.3%) N=1752	-776.9278 (146.1038) [0.0000] -0.125210 (0.028578) [0.0000]	Momentum <4.06bp
	-3000< $r_{-\tau}$ ≤-1000 (-0.3% to -0.1%) N=5412	-16.19772 (101.3537) [0.8730] -0.046879 (0.054829) [0.3926]	Uncertain Possible reversion 0.26bp to 1.20bp
	$\lvert r_{-\tau} \rvert$<1000 (less than 0.1%) N=15183	61.09890 (15.82826) [0.0001] 0.034101 (0.029581) [0.2490]	Likely momentum <0.90bp to Reversion 0.22bp
	1000≤ $r_{-\tau}$ <3000 (0.1% to 0.3%) N=5784	278.3712 (93.60108) [0.0030] -0.230926 (0.050502) [0.0000]	Momentum 0.42bp to Reversion 4.19bp
	$r_{-\tau}$ ≥3000 (more than 0.3%) N=1789	-139.5605 (155.5779) [0.3698] 0.000938 (0.032865) [0.9772]	Uncertain Possible reversion 1.42bp

Table 5: Regression results for price response from lagged price move for USD-CHF currency pair separated by size of price move for a timescale where the market maker likely does not resolve inventory shocks (20 mins.) and a time scale where he has likely been able to resolve his inventory imbalance (4 hrs). The 4 hour price responses (right column) are de-trended by -0.05bp.

Per the model, trend-chasing after an initial price move requires that arbitrageurs have a low level of risk-aversion (relative to the MM) because prices are further away from fair value. It does not seem reasonable that such arbitrageurs exist in significant proportion. Therefore, it seems likely that arbitrageurs generally trade counter-trend. Counter-trend trading and the MM's inventory control via the inter-dealer market cause small price moves to revert. The arrival of noise trader order flow which pushes prices further away from fair value leads collaterally constrained arbitrageurs to reduce (or reverse) their order flow. The MM then posts an even more deviated price—a momentum effect.

In the empirical analysis, there seems to be strong indication of small move reversions and some indication of big move momentum at high frequency. At a lower frequency, small move reversal exists along with some reversion and/or attenuation of earlier momentum.
The difficulty in finding clear reversion for big moves may be due to a variance in the timeframe for the MM to rebalance inventory. A variety of factors, including changes in market collateral and composition, affect the formation, growth, and deflation of micro-bubbles. Measuring across information events may also confound the analysis. Sporadic large jumps distort the measurement of the effect of small changes. This issue may be addressed with a longer duration dataset. Furthermore, it is likely that information events do not occur as "jumps" due to non-instantaneous dispersion of information. This phenomenon, along with traders chasing insiders, would generate longer stretches of momentum and more volatility and noise.

3.3   ALTERNATIVE EXPLANATIONS

Alternative explanations ought to account for both small move reversion and large move momentum. High leverage by itself is not sufficient to explain momentum.

If traders' actions are both synchronized and correlated with past prices then the market is not weak form efficient as it is possible to predict order flow and thus prices. As there is some statistical regularity in observed FX prices, it seems reasonable to assume that some traders' orders are synchronized—a significant portion of traders have taken similar positions. Schulmeister (2006) shows that commonly used FX technical trading strategies tend to produce signals on the same side of the trade and have substantial impact on order flow. The conclusion that traders may have synchronized trades and that prices exhibit momentum is consistent with some sort of trend-following trading strategy as well as the explanation suggested in this paper. Disciplined technical traders may close a position when they see a "trend reversal" or some other synchronizing signal. Of course, if these technical strategies are known in the market, it is likely the "reversal" would occur even earlier and the trend may break down before it even begins. It may be, then, that momentum effects arise from the market not fully knowing which strategies are being used and that the observed inefficiencies are the cost of learning.

Momentum may occur from a slow diffusion of public information and/or under-reaction as in Hong and Stein (1999). It seems not very likely that information events occur at a high enough frequency to generate the micro-bubbles observed. However, if the timing of information events is uncertain, trend-chasing may then also be a response to the perception of insiders trading— that is, traders following trades expected to come from private information. If true information events are relatively rare, it seems this type of trend-chasing would be arbitraged away. Thus, this alternate explanation seems to imply a high frequency of information events and insider trading.

The model also predicts that dealers hold "shadow" inventory from arbitrageurs. Empirical verification of this prediction may also suggest for the model over alternative explanations. As data on dealer inventory was not available, this test was not possible.

4 CONCLUSION

The model presented here builds on a microstructure approach to the FX market and ascribes price movements to order flow and a risk-averse market maker. Arbitrageurs, who trade counter-trend, appear to risk-share with the market maker. This trade actually extracts surplus from the market maker if an arbitrageur can wait for price to revert and then close his position. The market maker reacts strategically to this behavior by shifting price in the direction of arbitrageur inventory. The market maker also maintains a shadow inventory of arbitrageur positions. Collateral constraints on arbitrageurs are crucial to generating this outcome.

These price shifts as arbitrageurs accumulate inventory are a momentum effect which eventually reverts. Limited analysis of FX price data seems to provide some evidence of intra-day micro-bubbles.

The analysis presented here deals with arbitrageurs in aggregate. Individual arbitrageurs may be able to outperform substantially by engaging in momentum trading at appropriate times such as when a large amount of arbitrageur collateral has been committed to a counter-trend trade. Determining these events may be difficult from public information alone, though not impossible as the success of certain technical trading strategies has shown. Of course, a trader who has access to non-public information on market leverage and positions could do much better.

5 REFERENCES

Abreu, Dilip, and Markus K. Brunnermeier. (2002) "Synchronization risk and delayed arbitrage," *Journal of Financial Economics,* vol. 66(2), pp. 341-360.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. (1990a) "Noise Trader Risk in Financial Markets," *Journal of Political Economy,* vol. 98(4), pp. 703–38.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. (1990b) "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *The Journal of Finance*, vol. 45(2), pp. 379–395.

Evans, Martin D. D. and Richard K. Lyons. (2002) "Order Flow And Exchange Rate Dynamics," *Journal of Political Economy*, vol. 110(1), pp. 170-180.

Feng, Lei, and Mark S. Seasholes. (2005) "Do investor sophistication and trading experience eliminate behavioral biases in financial markets?" *Review of Finance,* vol. 9(3), pp. 305-351.

Garleanu, Nicolae, and Lasse H. Pedersen. (2011) "Margin-based asset pricing and deviations from the law of one price." *Review of Financial Studies*, vol. 24(6), pp. 1980-2022.

Gromb, Denis and Dimitri Vayanos. (2002) "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of Financial Economics*, vol. 66(2), pp. 361-407.

Hong, Harrison, and Jeremy C. Stein. (1999) "A unified theory of underreaction, momentum trading, and overreaction in asset markets." *The Journal of Finance,* vol. 54(6), pp. 2143-2184.

Kondor, Peter. (2009) "Risk in dynamic arbitrage: the price effects of convergence trading." *The Journal of Finance,* vol. 64(2), pp. 631-655.

Kyle, Albert. S. (1985) "Continuous auctions and insider trading." *Econometrica,* vol. 53(6), pp. 1315-1335.

Locke, Peter R., and Steven C. Mann. (2005) "Professional trader discipline and trade disposition," *Journal of Financial Economics*, vol. 76(2), pp. 401-444.

Lyons, Richard K. (1997) "A simultaneous trade model of the foreign exchange hot potato," *Journal of International Economics*, vol. 42(3-4), pp. 275-298.

Lyons, Richard K. (2001) *The microstructure approach to exchange rates*, Cambridge: MIT Press.

Mendel, Brock, and Andrei Shleifer. (2012) "Chasing Noise," *Journal of Financial Economics*, vol. 104(2), pp. 303-320.

Menkhoff, Lukas, and Ulrich Schmidt. (2005) "The use of trading strategies by fund managers: some first survey evidence," *Applied Economics,* vol. 37(15), pp. 1719-1730.

Menkhoff, Lukas, and Mark P. Taylor. (2007) "The obstinate passion of foreign exchange professionals: technical analysis," *Journal of Economic Literature*, vol. 45(4), pp. 936-972.

Odean, Terrance. (1998) "Are investors reluctant to realize their losses?" *The Journal of Finance,* vol. 53(5), pp. 1775-1798.

Osler, Carol L. (2003) "Currency orders and exchange rate dynamics: an explanation for the predictive success of technical analysis," *The Journal of Finance,* vol. 58(5), pp. 1791-1820.

Osler, Carol L., Alexander Mende, and Lukas Menkhoff. (2011) "Price discovery in currency markets," *Journal of International Money and Finance*, vol. 30(8), pp. 1696-1718.

Roche, Hervé, Stathis Tompaidis, and Chunyu Yang. (2013) "Why does junior put all his eggs in one basket? A potential rational explanation for holding concentrated portfolios," *Journal of Financial Economics*, vol. 109(3), pp. 775-796.

Schulmeister, Stephan. (2006) "The interaction between technical currency trading and exchange rate fluctuations," *Finance Research Letters,* vol. 3(3), pp. 212-233.

Shefrin, Hersh, and Meir Statman. (1985) "The disposition to sell winners too early and ride losers too long: Theory and evidence," *The Journal of Finance*, vol. 40(3), pp. 777-790.

Shleifer, Andrei, and Robert W. Vishny. (1997) "The limits of arbitrage," *The Journal of Finance*, vol. 52(1), pp. 35-55.