

Trading the currency markets: Is gut feeling still state-of-the-art?

By Richard Olsen, Olsen Group¹

The first shall be last, and the last first.

In the 1990s every large bank that wanted to be taken seriously placed a huge bet on their trading engines. But it was a long shot: they hired rocket scientists with the mandate to develop trading models based on neural networks, genetics algorithms and other magical tools that had paid off in the more respected disciplines of hard science.

The payoff in finance? We're still waiting. After a euphoric beginning the bubble of promise burst when live trading began and sudden profits failed to materialize. The projects were discreetly buried and the rocket scientists were either put on the street or assigned to other responsibilities.

Ironically, almost everyone at those banks was grateful for the failure. Because old-fashioned trading based on traditional experience and gut feeling suddenly got a new lease on life. And the lucky few who got the trends right could rake in large profits.

What went wrong? Why did the modern armory of sophisticated financial and computer engineering fail?

¹ Edited by Thomas Grizzard

These projects were begun in the hope of finding The Holy Grail. Program advocates at the banks were driven by the hope of rapid profits; enamored of modern computing technology, they expected the problem to be easy—plug and play. And because they were loath to share the spoils of the expected victory they failed to do what practitioners in the natural sciences have always done: invest time in thoroughly researching the data; observe the research object in every detail; run extensive tests before articulating a new theory; and then validate the theory with even more tests. Such classical rigor was deemed unnecessary. What the banks wanted was a magical solution. And they wanted it fast.

The legendary failure of the early attempts at trading magic is a sad testament to ambition. And to a lack of understanding. Financial markets are complex, dynamic and abstract—all of which makes the researcher's job especially challenging. And all of which invokes precisely the sort of rigorous investigation and testing the banks chose to bypass.

A small group of economists has bucked the trend and chosen the longer path of natural science. They collected high-frequency market data and analyzed its statistical properties. And to their surprise, the behavior of market prices contradicted the expectations of classical economic theory. Price distributions did not follow the expected Gaussian bell but were fat-tailed—with dramatic variations four times greater than “what they should have been.” And for extremely short intervals of time price extremes were even larger: up to 10 times the predictions of the classical model.

Why? What was driving this unexpected deviation? Well, it was no data fluke. Given the time extent of their analytical base and the millions of

transactions it captured, there could be no doubt: *actual markets behave differently from what classical theory had postulated.*

Nice to know, but not enough. So the renegade economists analyzed the detailed functioning of a financial market and traced *the flow of money*:

Whenever a trade takes place, the speculator or portfolio manager sells his asset via an intermediary to a market maker.

The market maker keeps the asset on his book and waits until another customer wants to do the opposite trade

or

The market maker goes to a fellow-market maker and hedges his position.

In the process, the market maker skews his price quote: if a customer has been *selling* him USD, he will skew the price of USD *down* to deter other market participants from selling him additional USD; if his customer has been *buying* USD, he will skew the price *up* to deter other participants from buying even more USD.

This skewing of prices—an essential strategy for the market maker to manage his risk—is a signal mechanism whereby the market maker can direct the flow of buying and selling. Tick-by-tick price movements reveal these buy and sell flows. Which means that high-frequency data are not meaningless noise but a detailed and information-rich footprint of the flows in the market.

By analyzing price changes tick by tick, researchers can accurately map the flows of a given market. Taken in isolation this information is of little value, since buy and sell flows alternate rapidly and price action is highly variable. And variability is a function of the continuous change in market liquidity. To understand liquidity, you have to understand market participation.

More dramatically for currency trading than for other assets, a key determinant of liquidity is the changing *presence* of market participants. In the global, we-never-close foreign exchange markets there are three distinct groups of traders: Asians, Europeans and Americans. These three groups obviously have different working hours but also unique preferences and institutional constraints that affect their trading behavior. And even more potentially dramatic is each group's characteristic response to news shocks and pricing action.

Researchers have traditionally analyzed the responses of traders *in physical time*—an approach suitable for inanimate objects, but inappropriate for human beings with brittle nervous systems. Such an approach does not adequately reflect the subjective experience of time.

High-frequency economic models refine the analysis by introducing the notion of *intrinsic time*. Intrinsic time weights chronological flow according to price action: during highly volatile periods time is expanded; during quiet periods, when market volatility is low, time is compressed. And this in fact mirrors our actual experience in the real world, where intense moments seem to pass much more slowly and quiet periods fly past.

The real payoff

By combining the analysis of buy and sell flows and the global time-zone effect, and applying the concept of intrinsic time, *it is possible to make critical and relevant observations about price movements*. Analogous to large-scale weather-forecasting systems, high-frequency finance models process market information and generate a weather forecast for the financial markets.

Much more complex than anything developed in the past, this approach provides valuable inputs for quantitative trading strategies, currency overlay, risk management and systemic risk analysis by central banks.

Compared to the highly sophisticated systems developed in the natural sciences, high-frequency analytical models are mere prototypes that prove the validity of the approach. But its ultimate utility remains to be seen. Just as biotechnology was scoffed at by the controlling pharmaceutical firms in the early 1980s, many experts view high-frequency finance as simply another quest for The Holy Grail.

We should be cautious but not timid: if we have learned anything over the last two decades it is that the pace of change exceeds our expectation and, often, our comprehension. The next big thing in finance may depend on nothing more than one spectacular success.

[1,100 words]