Portfolio Balance, Price Impact, and Secret Intervention*

Martin D.D. Evans and Richard K. Lyons

Portfolio-balance models occupy an important place within exchange rate economics. They are still used, for example, as a basis for why central bank intervention in foreign exchange markets can be effective even when sterilized (i.e., even when intervention has no effect on interest rates or money supplies). Empirical evidence supporting these models is scant, however. This paper tests the portfolio-balance model in a new, more powerful way and finds it strongly supported.

Past work on portfolio balance across assets denominated in different currencies falls into two groups: (i) tests using measures of asset supplies; and (ii) tests using measures of central bank asset demand. Here, we address the demand side, but we examine demand by the public broadly, rather than focusing on demand only by central banks. Under floating rates, changing public demand has no direct effect on interest rates, current or future. This provides an opportunity to test for whether price effects play a role in achieving portfolio balance. Because data on public trades became available only recently (with the advent of electronic trading), this strategy is feasible for the first time.

The discriminating power of our approach arises from avoiding difficulties inherent in past approaches. The asset-supply approach, for example, has low power because measuring supplies and their variation over time is

---

* The complete version of this paper can be found at:
< www.haas.berkeley.edu/~lyons/wp.html >.
Also published as National Bureau of Economic Research Working Paper No. 8356.

notoriously difficult. First, one must determine which measure of supply is the most appropriate. (There is considerable debate in the literature about this issue.) Then, for any given measure, the consistency of data across countries is a concern. Finally, these data are available only at lower frequencies (e.g., quarterly or monthly) and are rather slow-moving, making it difficult to separate the effects of changing supply from the many other forces moving exchange rates.

The central bank demand approach—an "event study" approach—may also have limited statistical power because central bank trades in major markets are relatively few and are small relative to public trading. For example, the size of the average U.S. intervention in the 1980s was only about $200 million, or roughly one-tenth of one per cent of the daily spot volume in either of the two largest markets. (In the 1990s, U.S. intervention was more infrequent but larger, typically in the $300 million to $1.5 billion range; at the same time, market volume was higher too.) Though the central bank demand approach is more successful than the supply approach in finding portfolio-balance effects, results are not exclusively positive and the extent to which these event studies pertain to portfolio effects more broadly is not clear.

The "micro portfolio-balance model" we develop embeds features more familiar to models from a sub-field of finance called microstructure. For example, the model clarifies the role played by a variable called "order flow" in conveying information about shifts in traders' asset demands. (The concept of order flow comes from microstructure finance and refers to *signed* volume. For example, a sale of 10 units by a trader acting on a dealer's quotes is order flow of –10.) Beyond clarification of the basic role played by order flow, two analytical results in particular are important guideposts for our empirical analysis: (i) order flow's effect on price is persistent (even when beliefs about future interest rates are held constant); and (ii) when central bank trades are sterilized, conducted secretly, and convey no policy signal, then the price impact of these trades is indistinguishable from that of public trades. This latter result links our analysis directly to intervention operations of this type.

With the advent of electronic trading and the data it provides, we have a powerful means of testing our portfolio-balance model. We establish three main results. First, testable implications of our model are borne out: we find strong evidence of price effects from portfolio balance. Thus, the portfolio-balance approach—with its rich past, but lack of recent attention—appears to warrant some fresh consideration. We consider this "resurrection" of the portfolio-balance approach the paper's most substantive contribution,

versus, for example, the paper's more narrow implications for intervention policy.

Our second main result is the precise estimate we provide of the immediate price impact of trades: about 0.5 per cent per $1 billion (of which about 80 per cent persists indefinitely). With gross flows in the largest spot markets at about $300 billion per day, this level of price impact is potentially quite important. An immediate example of this fact's value is its ability to help us understand why portfolio-balance effects from sterilized intervention are so hard to detect: the average-sized intervention of $200 million in the 1980s translates into an exchange rate movement of only 0.10 per cent, an amount easily swamped by movements due to other factors.

Our third result speaks to intervention policy. It clarifies how the unconditional price impact of 0.5 per cent per $1 billion varies with the state of the market. The most important state variable for the size of this price impact is the flow of macroeconomic announcements. (It may be, for example, that order flow is the variable that market participants use to resolve uncertainty about how these announcements are interpreted.) Whatever the reason, our estimates imply that trades have the most price impact when the flow of macroeconomic news is strong.

When applying our price-impact results to intervention policy, it is important to keep in mind that intervention in practice rarely takes the form we are examining in our paper (i.e., all the trades in our model are anonymous). Our analysis simply cannot address intervention that is conducted in any way other than anonymously. We realize that there is something of a consensus among economists that transparency (i.e., non-anonymous trading) is good when it comes to intervention. Our paper is an attempt to shed light on the other pole of the transparency spectrum—where intervention trades look exactly like anonymous private trades. This is a well-posed question, even if implementation of secret intervention is more challenging. If the central bank can trade anonymously, then one should expect its trades to match the average anonymous-trade price impact in the market. This conclusion does not require one to take a position on what is driving the rest of that order flow in the market.

Our paper's theoretical model provides an additional insight into why past sterilized intervention has often been ineffective. In our model, if order flow is expected to be reversed, it will not have much price impact. When customer position changes are permanent, order-flow effects on price have lasting portfolio-balance effects. In practice, after intervening, many (most?) central banks engage in systematic rebalancing of their currency positions (i.e., they tend to return to their pre-intervention portfolio composition).

This is problematic from a theoretical perspective if the objective is lasting exchange rate effects.

Finally, we offer some thoughts on future application of our trading-theoretic approach to intervention. Market data now coming available allow for precise tracking of how the market absorbs actual central bank trades and any information in them. Central banks with precise knowledge of their own trades—announcements, timing, stealth level, etc.—can estimate the impact of these various "parameter" settings. Consider the fact that the type of data used in recent work includes the order books of electronic interdealer brokers. A central bank with these data, over a sufficiently large number of intervention trades, can learn exactly how the "book" is affected, including the process of price adjustment, liquidity provision on both sides, and transaction activity. The situation (process) brings to mind a doctor who has a patient ingest blue dye to determine how it passes through the system. The whole process becomes transparent. Such is the future of empirical work on this topic.