

3

Psychology and Decision Making in the Foreign Exchange Market¹

Printed May 23, 2008

8872 words

Here begins the process of building the Post Keynesian alternative, something that will continue through chapter five. By then, a formal model will have been developed whose features include equilibrium trade imbalances, less-than-full employment, endogenous money, exchange rates marked by volatility and bandwagon effects, and an explanation of market participants' forecasts wherein expectations are guided by a mental model that is molded by social forces. The goal of this chapter is to examine the institutional structure of global currency markets and the social and psychological factors affecting market participants' forecasting and decision making. Among the phenomena explained will be volatility, bandwagon effects, and forecast-construction bias. Reference will be made to Institutionalism, Keynes' *General Theory*, and the psychological research of Kahneman and Tversky.²

THE INSTITUTION OF FOREIGN CURRENCY TRADING

The first step in understanding the currency markets must be a review of its organization. Agents

participating in the market for foreign exchange enter at one of three levels: wholesale (or market making), retail, or commercial. These are illustrated in Figure 3.1. Those acting in the top box are willing to make two-way offers on a continuous basis, such that they stand ready to buy *and* sell the currencies in which they offer wholesaling services. Agents operating at the wholesale or retail level may contact a wholesaler and request prices for a currency. Both a buy and sell rate will be quoted by the wholesaler before the caller reveals their intention. The caller then decides whether or not they will carry out their transaction at the price offered. The goal of the wholesaler is to “generate revenues from the *spread* between the offer and the bid” (Shoup 1998: 105), or the difference between the buy and sell prices. As spreads at this level will be very small, wholesaling depends on generating volume to create income. Banks and other large financial institutions capable of maintaining well-staffed and equipped currency trading rooms with the liquidity necessary to maintain large inventories are typically the only ones willing to offer this service.

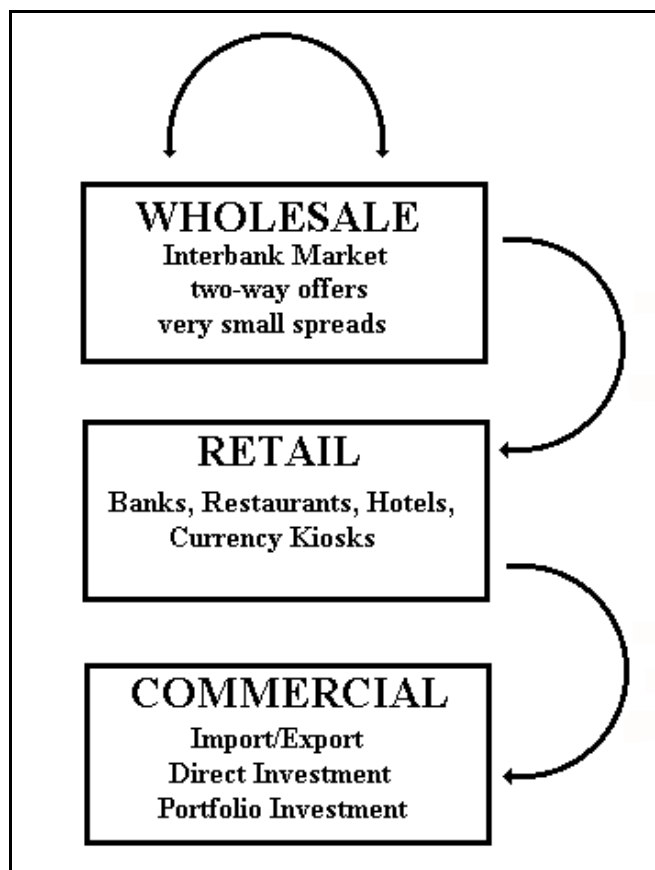


Figure 3.1: Foreign Currency Market Structure.

Retailing involves making one way offers. Within the context of a particular transaction, those so acting are willing only to buy *or* sell a particular money. They purchase their funds from an agent who is wholesaling (who will, as indicated on Figure 3.1, offer two way prices to the retailer) and then generate their income by selling at a markup to their customers (as indicated by the link between Retail and Commercial in Figure 3.1). Branch banks, restaurants, and hotels commonly undertake retailing.^{3,4}

Those acting commercially are importers, direct investors, and portfolio investors.⁵ Importers buy (or contract to buy) foreign currency because they wish to purchase foreign goods or services.⁶ Direct investment is long-term capital investment abroad, such as the establishment of a multinational subsidiary or the purchase of a significant interest in some foreign firm. Portfolio-capital investors purchase the currencies necessary for the acquisition of the international financial assets (including deposits of the money itself and any associated interest income) they want to add to their portfolios. It is assumed that such purchases involve no long-term commitment on the part of the buyer and have as their goal short-run capital gain.

A single agent could undertake activities at all three levels (wholesale, retail, and commercial) at various times or in various markets or currencies. A financial institution could, for instance, operate in the wholesale market at one moment, supply funds to customers at the retail level in another, and engage in portfolio investment activities as a commercial actor in yet another. Playing multiple roles has the potential to afford such agents a competitive advantage since they can internalize some of the costs. That large institutions might be best placed to do this means that

there is a potential for concentration in the market for foreign currency. Indeed, the 2002 Bank for International Settlements report on foreign exchange activity has reported trends consistent with this possibility:

The consolidation trend in the banking industry that started in the mid-1990s appears to have continued between 1998 and 2001...In the United States, 75% of forex market transactions were conducted by only 13 banks in 2001 compared to 20 banks in 1998 and about 20 banks in 1995. In the United Kingdom, 17 banks captured 75% of the market in 2001 compared to 24 banks in 1998 and about 20 banks in 1995.

(BIS 2002: 9)

There is evidence that market participants sense this well, as shown in Cheung and Chinn (2000). Their survey of currency dealers indicates that for the big four currency rates, 17% thought that there were dominant players in the dollar-Deutsche Mark market, 22% in the dollar-yen, 50% on the dollar-sterling, and over 58% in the dollar-Swiss franc. The reasons survey respondents supposed that such agents could dominate were related primarily to sheer size.

CURRENCY MARKET PARTICIPANTS' ROLES

Prices in the foreign exchange market are negotiated among market participants on a continuous basis. As few rigidities or externalities exist, the question of what determines the exchange rate

becomes what determines the relative demands for currencies? To answer this requires a closer look at wholesaling, retailing, and commercial demands.

The effect of retailing is relatively minor. Retailers merely intermediate between commercial actors and wholesalers, charging a mark up in the process. This means that retailing has no independent impact on the currency market. They act in response to their customers' demand and earn income from the markup they charge.

The act of wholesaling is more complicated. Because they earn their income on the narrow spread between bid and ask prices, they must quote prices that they *anticipate* will that generate an equal volume of business on both sides of the market. If they are incorrect, then their inventories of currencies accumulate in a way that creates for them a vested interest in future exchange rate movements. For example, if a currency-trading desk set a price that attracted more sell orders for the yen than buy orders, it would accumulate large inventories of that currency. In that case, they would rather not witness a yen depreciation! Wholesalers by definition want to earn risk-free profits by "jobbing." It is not their goal to engage in currency speculation and thus quoting a price that will leave working inventories unaffected is a central goal.

Doing so is not simple. As the currency desk receives orders over the course of the day, the wholesaler must decide whether the timing of the orders is coincident or the result of unexpected trends (Suvanto 1993: 1-22). If it is the former and, in fact, the price the wholesaler is quoting will (by the close of business) leave her with a closed position (i.e., without unintended inventory

accumulation), then the fact that the morning happened to be witness to an excess of orders for the euro, for example, should not lead her to change her bid or offer. But, if it is determined that the unexpected demand for the euro is a function of an actual shift in market sentiment then the price at which the euro is sold must be increased if desired inventories are to be maintained. Thus, it is incumbent upon wholesalers to make careful forecasts of customer demands (something that is not necessary in the more passive act of retailing). That greater sophistication is required to wholesale is the reason that large banks and investment firms are usually the only agents in a position to do so (they will also almost certainly be among those firms with oligopoly power referenced above).⁷ Despite this greater complication, however, wholesaling has roughly the same impact on prices as retailing even if more steps may be involved. They are only acting in response to changes that have taken place (or that they anticipate to take place) at the commercial level. It is, therefore, from the latter that the ultimate demand for currency arises, and thus where exchange rates are determined.⁸

As suggested above, there are three basic sets of commercial demands for foreign currency: those arising from imports, from the demand for direct foreign investment abroad, and from the demand for portfolio investment abroad. A rise in any one of those will increase the demand for and (ceteris paribus) the price of foreign currency. When commercial actors wish to undertake any one of the three, they will contact retailing agents, who in turn place orders with firms who are wholesaling (again, this could sometimes be all within the same company). Barring offsetting demands, this will cause wholesaling inventories of the currency in question will be run down (while that of the currency be supplied by customers will rise). As they do not want to take a

position in the market, this may lead those wholesaling to compensate by raising the price of the currency in demand (thus lowering that of the surplus one; this may require several hedging operations before wholesaling agents are satisfied with their exposure). Ideally, of course, wholesaling agents will have anticipated this turn of events and raised the price before the first order arrived (thus keeping the flows on both sides of the market even and reducing the necessity for covering transactions). But in either event, currency prices move. Note that this is the case even with forward contracts as wholesalers will buy the spot equivalent immediately, selling it to the customer on the day the contract matures.

What drives commercial activity and hence exchange rates? Beginning with imports, these rise as domestic national income rises (and domestic agents wish to buy more goods and services, including foreign ones) and falls as the relative price of foreign goods and services rises.

Accordingly, a rise in US national income, for example, would lead to a rise in US imports and a subsequent dollar depreciation as the level of world demand for foreign currency (as compared to that for the dollar) would have risen. An increase in the relative price of foreign goods and services would do the opposite and hence cause a dollar appreciation.

There is substantial evidence, incidentally, that trade flows tend to be price inelastic but income elastic, such that large changes in exchange rates or national price indices have only a small and delayed effect on trade flows while fluctuations in overall economic activity (GDP, for example) has a quick and substantial impact (see, for example, Chinn 2005). We have had dramatic examples of large current account imbalances being very resilient in the face of substantial

exchange rate depreciations over the past several decades. Imports do respond to price changes, but the impact tends to be delayed and muted.

Direct foreign investment is somewhat more complex, but a pattern emerges nevertheless. To begin, direct foreign investment can either be vertical (along the stages of production of a given product; for example, a tire manufacturer may acquire foreign rubber plantations abroad) or horizontal (at the same stage of production; for example, if that same tire manufacturer establishes another tire manufacturing plant in a different country). Horizontal direct foreign investment is often market-seeking and therefore tends to be attracted to nations with high levels of income and similar tastes to those in the source country; vertical is typically resource-seeking and moves to nations with cheap sources of those resources (including labor) necessary to add value to the product in question. As direct foreign investment is attracted to an economy for whatever reason (e.g., rise in domestic income attracting horizontal or a fall in domestic prices attracting vertical), so its currency appreciates and that of the source country depreciates. Hence, a rise in national income might tend to attract horizontal direct foreign investment and lead to a domestic currency appreciation, while a fall in factor costs would do the same but via vertical direct foreign investment.

As explained in chapter one, however, as large and important as trade flows and direct foreign investment are, it is portfolio capital that dominates the market for foreign exchange. The factors driving financial flows are many but can in general be reduced at this stage to asset yield, default risk, and liquidity.⁹ The greater the yield an agent expects to earn, the greater the demand for the

asset; the greater the default risk, the less enthusiastic an agent is to buy; and the easier one suspects it will be to liquidate the asset, the more attractive it will be. Nations whose assets are perceived to be offering a higher yield, lower chance of default, and greater liquidity will experience appreciating currencies as agents rush to buy those assets, creating net capital inflows.

In considering the purchase of foreign assets, market participants must evaluate factors related to both the issuer of the asset in question (public or private) and those related to the currency in which the asset is denominated (relative to that used by the participant). Consider equation 3.1, showing the US dollar value of a sample portfolio of sterling denominated assets:

$$V = (\$/\pounds)*BND + (\$/\pounds)*DEP + (\$/\pounds)*P_{stk} *STK \quad 3.1$$

where V is the total US dollar value of the portfolio, ($\$/\pounds$) is dollars per pound sterling, BND is the sterling value of bonds held, DEP is the value of bank deposits in sterling, STK is the volume of individual stocks owned, and P_{stk} is their average price (in sterling). Note that, save any adjustments to BND and DEP that must be made to account for accumulated interest income or price adjustments caused by interest rate movements, the primary factor affecting the dollar value of the first two asset types is the rate of exchange. In addition, the relatively high volatility of ($\$/\pounds$) (especially as the agent diversifies the group represented by STK) means that it is at least as important as P_{stk} in the last expression, and probably more so. The bottom line is that when investors are forecasting foreign asset values, *expected changes in currency prices are a key factor*. Therefore nations with currencies that are expected to appreciate will attract portfolio capital flows, causing an immediate appreciation.

To say that this is a self-fulfilling prophecy at work is true at a very basic level, but it glosses over the fact that there may be a very reasonable and well-considered rationale underlying the prophecy (as will be argued in chapter five in the context of the mental model). Note first that the variables upon which agents must focus in generating these foreign exchange predictions depend on which activity (imports, direct foreign investment, or portfolio foreign investment) that they perceive as dominating the market. Were it, for example, that they believed that imports constituted the majority of foreign exchange transactions, then portfolio capital investors would be well served to pay careful attention to the determinants of the sales of goods and services across national borders (i.e., national income and relative prices) and react accordingly. For instance, if portfolio investors expected a nation's income to rise, then they would come to the conclusion that the nation's imports were about to increase. As this would eventually drive down the value of the currency and thus hurt the value of assets issued by that country, market participants would immediately sell those assets. The outflow of portfolio investment would then cause the very depreciation agents feared.¹⁰ This is a self-fulfilling prophecy, but one based on a reasonable evaluation of the circumstances. If the anticipated increase in that nation's income does not come to pass, then agents may revise their forecast and move the currency price in a different direction next period—but in the meantime, currency prices did change and economic activity was forced to adjust to the new level.

In point of fact, of course, it is portfolio investment and not trade that dominates world business, and in particular flows in search of short-term capital gain (Krause 1991, Schulmeister 1987, Shelton 1994, and Walter 1991). In that case, what agents monitor most are the determinants of

portfolio foreign investment. If the flows into a particular country are expected to increase then that nation's currency can be expected to appreciate; and if its currency can be expected to appreciate, then it is safe to assume that this will, *ceteris paribus*, attract portfolio capital. The expected appreciation subsequently causes actual inflows of financial capital, which leads to the spot market currency appreciation that was expected.

If this is indeed the manner in which the market operates then it is apparent that determining the process by which agents form expectations is vital to explaining exchange rates. It is also important to note that the foreign exchange forecast drives *today's* rate, not the one it is forecasting. Today's expectation of next week's spot rate affects today's spot rate. Whether or not the forecast turns out to be an accurate predictor of the spot rate is an interesting question, but it is not the central one. Our concern with the forecast is a function of its role as a driver of current rates, not as an indicator of future ones.

THE CHARACTER OF EXPECTATIONS IN THE FOREIGN CURRENCY MARKET

To reiterate, since portfolio capital flows dominate the foreign exchange market it can be said that *it is today's expectations of future currency price movements that play the most important role in determining the current foreign exchange rate.* Today's prices are created by the weighted (by liquidity and confidence) average of market participants' expectations of tomorrow's price. Agents are not, as in rational expectations, forecasting an event that is independent of their actions—they are creating the event (Davidson 1982-3). Realized outcomes clearly affect on the

exchange rate (as, for example, importers demand foreign currency and cause the latter to appreciate), but, even then, the current structure of the currency market means that they do so primarily through expectations. The dollar moves more in reaction to the *announcement* of a trade imbalance than from the pressures created by the imbalance itself.

Before considering how the expectations are formed and decisions made, note that if the above characterization of the currency market is accepted then the standard Neoclassical conception of forecast bias takes on less importance (while another one becomes more so). In Neoclassicism, a bias is a persistent forecast error. Forecasts are not assumed to be usually or even necessarily ever correct; but, if economic agents are rational, then any errors must be random. This is because systematic errors, by nature, can be identified and eliminated, something rational, profit-seeking agents would have a strong motivation to do.

In the world described by Neoclassicism, the lack of bias, so defined, is an important indicator of rationality and market efficiency, two of their core concepts. In their view, currency market participants' expectations, like those of weather forecasters, have no impact on the actual outcome. Back on June 4, what I and my colleagues *think* the euro will be worth on July 4 has no effect on what the euro is worth on July 4. And the amount by which my forecast ultimately misses the mark is not only a major factor in the value of my portfolio, but the pattern of such errors says something about me. If I (or all of us in aggregate) continually overestimate but I do not adjust my subsequent forecasts accordingly, then I am irrational. In addition, the fact that there are others who are rational and will make such adjustments means that they will drive me

out of the market. Hence aggregate market expectations are bound to be rational (at least in the long run) because either all the participants are individually rational or, at the very least, that portion of the market that is irrational will be driven out of business. Note that one of the strengths of this approach is the fact that how market participants form their forecasts is irrelevant. Little time is spent pondering that question in Neoclassicism. Rather, it is the difference between the forecast (however created) and the realized value that is central, and logically so given their view of the phenomenon.

But because from the Post Keynesian perspective it is the forecast that creates the realized values, understanding the process by which the former is created is vital. Meanwhile, by the time the realized value associated with a particular forecast is known, the original forecast is irrelevant as market participants' energies are by then focused elsewhere. Imagine, for example, that it is Monday morning and a representative agent has spent most of the day making a forecast of the position of the euro by the following Monday (for sake of argument, take a forecast time horizon of one week as being the standard). She eventually decides that it will appreciate relative to the dollar, rising from \$1.20 (the current spot rate) to \$1.25. For simplicity, say that this is the consensus view and she and all her colleagues world wide have acted on this expectation.

Further assume that their confidence in this forecast is absolute such that they continue buying the euro throughout the rest of the day until they, collectively, have driven the current price to \$1.25.

Tuesday dawns for our representative agent and, barring any change in expectations, the euro price still stands at \$1.25. This means that, assuming she acted before the price reached its

maximum, her portfolio has increased in value (which was, of course, the whole point of the exercise). Does she now stand pat and wait to see if yesterday's forecast comes to pass, at which point she empties her inventory of euros? Of course not. She plans on holding those euros until conditions suggest that she should sell at least a portion of them (or buy more). Today, she works on her forecast for next Tuesday. In particular, she wants to know if there is any reason to believe something other than what she thought yesterday. A great deal of time and effort will be devoted to this and as the day wears on, so she will develop a new one-week forecast and alter her portfolio based upon it. The aggregate impact of her colleagues all over the world doing the same thing will move the exchange rate. On Wednesday, the process starts over again.

Come the following Monday, will she stop to compare the day's realized price with her forecast from a week earlier? Odds are, she will not, for a couple of reasons. First, her day will be busy enough generating her next forecast. It is not immediately evident that there is more to be gained from analyzing the last week's events than there is from monitoring today's currency markets. If time is the ultimately scarce resource, reading the European Central Bank's latest statement (and market reactions thereto) is likely to be far more profitable than reliving the past seven days. Furthermore—and this is the more critical point—her forecast one week ago was made *ceteris paribus*. But everything else did not stay equal. Perhaps by today the euro really stands at \$1.23 (recall that her original forecast was \$1.25). This may cause no concern whatsoever on the part of our representative agent, because she may feel that intervening events fully justify this new level, rendering her original forecast moot. She is not terribly concerned with the difference between her forecast and the realized value seven days later; however, what would have caused her great

consternation would have been if one week ago when the euro stood at \$1.20 and she forecast it to move to \$1.25 (and therefore bought euros), the euro had fallen to \$1.19 by the next day.

Under those circumstances, *regardless of how accurate her forecast eventually proved to be*, the fact that her expectation was out of line with that of the majority of the market on day one will be a source of considerable worry. This is Keynes' beauty contest, and it is the critical problem for the representative agent. This is why the focus in the Post Keynesian world is the difference between today's forecast of next week and today's realized price—not next week's.

There cannot be, of course, any difference between today's aggregate forecast of next week's rate and today's realized price (save that created by the lack of forecast confidence that may lead agents to hedge their bets). This is a matter of concern only to individual agents. So, a) while bias in the Neoclassical sense can still exist, it is not a theoretically significant issue, and b) there simply cannot be a "bias" gap between today's average market forecast and today's realized price.

This does not mean that there is no concept of bias in the Post Keynesian perspective, however. As will be seen below, modern psychology argues that the factors guiding decision making may introduce of a number of varieties thereof. But in this context, "bias" does not mean forecast error; it is instead the unreasonable influence of some factor in the formation of the forecast.¹¹ This unreasonable influence, if widely shared, does not lead to a mistake. It instead becomes part of the realized price. It is important not because it indicates the rationality or irrationality of market participants, but because it influences the objective variable and economic activity is forced to adjust to it. Note that this bias (theoretically, if not by the agents) is known on the day of the

forecast; the Neoclassical one cannot be known until the date of the object of the forecast has come. To distinguish it from the latter (which will hereafter be referenced as realized-forecast bias), it will be called forecast-construction bias.

DECISION MAKING AND FORECASTING IN THE FOREIGN CURRENCY MARKET

Because Post Keynesian economists see agents' expectations as key determinants of foreign currency prices, explaining the latter is an essential step in building a theory of exchange rate determination. Doing so will require reference to the work not only of economists (primarily Keynes), but of psychologists as well. Considerable space will be devoted to this task as some of the concepts may be quite new to the reader. Once the basics have been described, it will be possible to explain six salient features of the foreign currency market: forecast-construction bias, price volatility, bandwagon effects, technical analysis, trading limits, and profit taking.

Stages of Decision Making

How do people choose? Psychologists argue that humans “rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations” (Tversky and Kahneman 1974: 1124). These heuristics enter into the decision-making process at various stages, with the latter defined as:

1. Eventuality Analysis;
2. Choice and Consequence Definition;
3. Decision Weight Assignment;

4. Choice;
5. Post-event Assessment (Harvey 1998).

In Eventuality Analysis, the actor considers all the probable future states of the world (as related to the decision to be made). Under Choice and Consequence Definition, each possible future (from step 1) is compared with the alternatives available to the actor and the consequences of the interactions of each future and choice are contemplated. The matrix below gives an example of this process of interacting possible futures with choices.

	euro appreciation	euro depreciation
buy euros	profit	loss
sell euros	loss	profit

For simplicity, imagine that there are only two possible futures: euro appreciation or depreciation.¹² Further assume that the agent has only two choices available: short the euro (sell it) or go long in the euro (buy it). Each cell shows the consequence of that future combined with the choice on the vertical axis. If the agent were to buy the euro and the euro appreciated, a profit would result. If they had gone long in the euro but the euro depreciated, that leads to a loss, and so on. The agent thus compares each possible future with an available choice and forecasts the consequences.

Decision weight assignment arranges the alternatives in order of preference (based on the analyses performed in stage two and the agents' estimates of the likelihood of each outcome). In terms of the choice/consequence figure above, this means ranking the alternatives "buy euros" and "sell euros." As a first approximation, this can be seen as equivalent to calculating expected values. Continuing with the above example, if the agent in question believed that there was a seventy per

cent chance of euro appreciation and a thirty per cent chance of euro depreciation, then one would expect the agent to rank “buy euros” first and “sell euros” second.

Stage four is the point at which the actor actually selects a course of action (basically, the one with the highest decision weight). Though this does not always occur, the decision maker may subsequently undertake Post-event Assessment, meaning that the choice and realized consequence are reviewed in the light of what was expected.

In both stages one and two, a major part of the process is the establishment of probabilities and confidence levels. More specifically, in stage one (eventuality analysis), one must decide both what might happen over the relevant time horizon and the relative likelihoods thereof, while in stage two (choice and consequence definition) one must project likely interactions (a range of potential outcomes and the probability of each) between the currently available choices and the possible futures from stage one. In the example above, these were rather straightforward. But, if eventualities had been a range of European Central Bank policies, there is no longer a clear correspondence between outcomes and profitable portfolio decisions. The effect of new financial regulations, for example, might not be obvious and thus each cell in the choice and consequence definition table above would contain multiple items, each with its own probability. This is likely to be very complex and most of our representative agent’s time will be involved in these stages one and two. Note that there must also be a level of confidence associated with every forecast.

Heuristics and Other Tendencies

Both the probability and confidence level determination that dominate eventuality analysis and choice and consequence definition require the use of the three main heuristics of human decision making: availability, representativeness, and anchoring. With availability, which is used to estimate frequency (in the past) or likelihood (in the future), the more available something is in memory (either through imagination or recalling past instances), the more frequent or likely that event is deemed. So, for example, because instances of snow in January are easier to recollect/imagine than instances of snow in August (in the Northern Hemisphere), agents will declare the former to be more likely than the latter. Such events come to mind more easily and are therefore believed to have a higher chance of occurring. Representativeness is used when one is concerned with the probability that object A belongs to class B (i.e., the chance that event A is the result of process B, or that process B will create event A). The more A resembles B, so the heuristic goes, the more likely that it belongs to class B. The series of coin tosses T-H-T-H-T may be deemed more likely the outcome of a random coin toss than T-T-T-T-H because the former better represents randomness (where the outcomes are event A and the tossing of a fairly-weighted coin is process B). Anchoring occurs when the individual must make a forecast. When this is done by starting at some initial estimate and then adjusting, people tend to anchor to that first value *regardless of the process used to generate it*.

All three of these are very useful under most circumstances, but each also introduces what I will call “forecast-construction bias.” For example, availability seems on the surface of it to be a very reasonable approach, and one that saves computational effort. Of course it is true that snow is more likely in January, and it is not necessary to consult any meteorological charts to come to this

conclusion. The inherent problem is, however, that *there are many things that can make something more available in memory without making it truly more frequent or likely*. An instance may be more easily accessed, for example, simply because it was dramatic, more recent, or falls into the decision maker's area of interest. The ease with which a scenario can be constructed or the search set used can also bias availability (which can be a particular issue when the agent is employing a specific mental model—phenomena that do not “fit” will be subconsciously ignored). This means that, in forecasting exchange rate movements, agents may overrate the importance of events that were more recent or dramatic or fit preconceptions. These are forecast-construction biases, as opposed to realized forecast biases, and they will be reflected in currency prices.

Like availability, representativeness is both a useful heuristic that allows us to arrive at reasonable answers with a minimum of effort and a method that carries with it inherent biases. In general, the issue with representativeness is that in considering whether or not outcome A is a function of some process B (or, similarly, what outcome A will be generated by process B), agents allow for very little variation between A and B. In other words, people expect a random coin toss to look like H-T-H-T-H-T, and are tolerant of only very little deviation from the fifty-fifty distribution of heads and tails. If a different pattern is evident, then agents assume that there must be at work some distinct process that creates precisely that event. Thus, people expect causation where chance may be dominant (thereby ignoring very simple and presumably intuitive rules of statistical inference). For example, an outstanding game by an athlete is seen to be a function of the athlete having been unusually skillful over that period. That the performance could be a result of a random variation from the mean is not considered.

If the law of averages is considered, representativeness means that it may be applied incorrectly. For example, agents typically expect that outliers in one direction will be offset by outliers in the other direction, even though there is no reason to suspect that tossing five heads in a row will cause the next five to display an above-average distribution of tails. That there are software programs, web sites, and books that purport to predict lottery numbers on the principle that numbers that have already been selected are less likely to be drawn a second time is attributable to representativeness. Each event is, of course, independent. Lottery numbers have no memory.

Representativeness also causes people to assume that they can make accurate predictions based on scanty information. In other words, knowing (or thinking they know) the nature of process B causes agents to believe that they can forecast outcome A. In the foreign exchange market, representativeness means that agents expect every currency movement to have a very specific cause, whether they can discern it or not. The effect of the timing of orders or other coincidences are interpreted as important market-driven trends. The willingness to forecast based on scanty information reinforces this tendency.

Anchoring puts undue weight on an individual's first estimate. Amos Tversky and Daniel Kahneman (1974) conducted a study in which individuals were asked to estimate the percentage of African countries in the United Nations. A wheel of fortune with numbers ranging from 1 to 100 was spun in front of them before they could answer, and once it had come to rest on a value, the subjects were asked to give their estimate plus or minus that value. Despite the fact that the

numbers were clearly generated randomly, agents anchored their responses to them! Likewise, because the most recent time series will inform any forecast, currency market participants will tend to anchor to levels in calm markets and rates of change in volatile ones.

Another probability-associated factor (though not a heuristic) relevant to stages one and two is that agents appear to believe that outcomes more beneficial to them are more likely.¹³ This is a common psychological tendency (sometimes referenced as “wishful thinking”). Hence, one can assume that profitable outcomes will, *ceteris paribus*, receive a greater weight.

Through all these processes framing is also a very important issue because

Peoples' probability judgments are not attached to events, but to descriptions of events (Tversky and Koehler 1994 and Tversky and Kahneman 1988). Events do not have likelihoods. A likelihood is an opinion, not an objective value. Opinion is based not only on the information available to the decision maker, but also on the framing of that information. Study after study has shown that the same question asked in different terms yields different answers, in stark contrast to the predictions of expected utility theory. Framing is an issue not only in the description of events, but also in the structure of the choice being made (Redelmeier and Tversky 1992). A change in either has the potential to change the decision that is made, and thus framing must be considered an important factor in the decision-making process and in probability assessment.

(Harvey 1998: 54)

Events must be interpreted and agents do this within a particular social context. They share beliefs about what causes outcomes, and this mental model is a major determinant of what they interpret as relevant and how they do so (this will be discussed further in chapter five).

The above—availability, representativeness, anchoring, wishful thinking, and framing—all deal with the manner in which probabilities are calculated in eventuality analysis and choice and consequence definition. The specific impact of these will be outlined later in the chapter. For now, suffice it to say that in the process of collecting information, forming forecasts, and making decisions agents will, when it is necessary to calculate a probability, generally rely upon the principles inherent in the above mechanisms and their mental model. Forecast-construction bias is a natural byproduct of the process; in particular, agents will place undue emphasis on dramatic and more recent events, ignore basic statistical principles, anchor to early estimates, have a tendency to expect events that favor them, and tend to ignore evidence that does not fit their preconceptions. These have specific consequences in the currency market.

Forecasts such as those formulated in stages one and two will also have associated with them confidence levels. The main principle is that the more easily the decision maker was able to make a probability judgment based on the available information (i.e., the greater the weight of the argument), the greater the confidence in that judgment. Also, if agents are engaging in Post-Event Assessment, then any forecast errors that they discover would reduce overall confidence. Conversely, the substantiation of earlier predictions increases confidence. Note that price volatility tends to, *ceteris paribus*, decrease confidence.

Stage three requires the assignment of decision weights to each alternative. As suggested above, at first glance this can be viewed as equivalent to the expected values of rational choice theory. There are a number of key differences, however. First, the agent will (*ceteris paribus*) prefer the option in which she has more confidence. If I determine that there is a 75% chance of a euro appreciation and that same probability of a yen appreciation, but I have more confidence in my yen forecast, then it will receive a higher decision weight. In addition, psychologists have discovered that people prefer risky options when they feel as if they are losing and safe ones when winning (Tversky and Kahneman 1992: 298). Agents holding appreciating currencies are therefore more inclined to select options that allow them to take profits, while those with depreciating monies will want to not only avoid realizing their losses, but undertake further deals that might create offsetting revenues. Last, agents will prefer the option that is most likely to allow them to claim credit and avoid blame. There exist, for example, strong incentives to follow the crowd, even the (*ceteris paribus*) an individual believes the best course lies elsewhere. Consider the following matrix showing the choices faced by a market participant (vertical) and the possible outcomes (horizontal):

	Correct Decision	Incorrect Decision
With Conventional Logic	Rational	Unlucky
Against Conventional Logic	Lucky	Ignorant

The four cells cross referencing choices and results show the manner in which the market participant's actions are likely to be interpreted (by peers, customers, supervisors, etc.). Note the

incentive to move with the crowd: someone going against the majority and losing is clearly an ignorant individual, while doing so and winning is only marginally better as it may be interpreted as luck; but those choosing to move with the crowd are rewarded with a positive characterization of their choice when correct, and they may receive at least a partial dispensation if incorrect.

Once decision weights are assigned, the choice stage is generally a matter of picking the alternative with the highest weight (unless agents lack sufficient confidence regarding a decision, in which case they may postpone it).

In stage five, agents *may* review past decisions with an eye toward refining one's decision-making. This can be useful, but it is often the case that the framing of decisions and means of data collection is not conducive to this process. To further compromise the matter, people tend to recall their past decisions as having been more successful than they actually were. Finally, as explained above, by the time the forecast date has arrived, agents may be focusing more energies on the next forecast than reviewing the past.

Keynes on Forecasting and Decision Making

Keynes' insights into the workings of asset markets add to the psychological view. He makes five observations that are particularly relevant here:

1. Uncertainty- Because in general our knowledge of the future is "vague and scanty" (Keynes 1964: 148), the information of which we are aware plays a disproportionate role

in our forecasts. For example, if event X depends on factors A, B, C, D, and E, but because of the nature of the world we are only able to know one piece of information at any given time (A, for example), it will play a larger role in our forecast of X than if A, B, C, D, and E were known.

2. Convention- Asset market participants adopt the convenient convention “that the existing market valuation, however arrived at, is uniquely correct in relation to our existing knowledge of the facts which will influence the yield of the investment, and that it will only change in proportion to changes in this knowledge; though, philosophically speaking it cannot be uniquely correct, since our existing knowledge does not provide a sufficient basis for a calculated mathematical expectation.” (Keynes 1964: 152).
3. Low Confidence- Since knowledge of the future is “vague and scanty,” levels of confidence are likely to be fairly low in asset markets (though this is offset by our animal spirits; see number five below). Forecasts are therefore “...liable to change violently as the result of a sudden fluctuation of opinion” (Keynes 1964: 154).
4. Quick Results- “...human nature desires quick results, there is a particular zest in making money quickly, and remoter gains are discounted by the average man at a very high rate” (Keynes 1964: 157).
5. Animal Spirits- Humans exhibit “animal spirits,” or a “spontaneous urge to action rather than inaction” (Keynes 1964: 161). This offsets any misgivings we may have about making important decisions in an uncertain world and allows us to act despite our ignorance.¹⁴

These will be melded into the psychological view in the analyses that follow.

SALIENT MARKET CHARACTERISTICS

Orthodox economics has struggled not just to explain overall currency price determination, but to account for commonly observed phenomena like bandwagons and bias. What explanations do exist tend to be inconsistent with parallel explanations in the same source (a common situation in textbooks) or relegated to the short run and, therefore, white-noise status. By contrast, the combination of the psychological view and Keynes' observations offered here yields a description of a decision making process that includes inherent forecast-construction biases, tends to create price volatility and bandwagons, and leads agents to employ technical analysis, take risks, and engage in periodic profit taking.

Forecast Construction Bias

That the first is true should be clear from the psychology. As has already been explained, persistent forecast-construction biases are associated with availability, representativeness, anchoring, and the tendency to expect favorable events and ignore those that do not fit preconceptions. These biases are incorporated into the price and thus those who try to "avoid" them do so at their own risk. Forecasting the psychology of market means treating these as important determinants and not errors.

Volatility

With respect to volatility, there is no question that it has been so great that it cannot be explained by changes in underlying “fundamentals” (Quinn and Harvey 1998). The Post Keynesian approach blames the volatility on the manner in which agents’ forecasts are formed. It is created by uncertainty, availability, representativeness, anchoring, the desire for quick results, animal spirits, and convention. Uncertainty means that as events emerge, so agents, lacking the whole picture, are likely to give them disproportionate weight and react accordingly. Representativeness works to reinforce this as it encourages people to think that they can and should make forecasts based on small samples. If the event that created the initial price movement was particularly dramatic then the availability heuristic leads agents to overweight it in their forecast; and once volatility starts, anchoring may exacerbate it as attention shifts from price levels to price changes (in other words, agents anchor forecasts to the changes such that they begin to expect them rather than stability). Throughout this process, agents’ desire for quick results and the willingness to act even when lacking a firm basis for decision making (i.e., animal spirits) makes the emergence of sudden and violent price movements, if not a daily event, certainly not surprising. One might think that convention—the belief that today’s price is uniquely correct given our knowledge of all that has transpired up to this moment—would encourage stability. It may under some circumstances. However, it also means that each time some new event occurs, agents feel bound to assume that further price changes are inevitably warranted. This may occur even when the “new” information is simply the continuation of an existing state of affairs (an interest-rate differential, for example).

Volatility feeds on volatility, but only to an extent. If it becomes too great, agents’ increasing lack of confidence (which is likely as prices become unstable) may eventually win out, causing them to

withdraw from the market (closing out their positions). Calm will then be restored, until events conspire to create a new period of volatility.¹⁵

Bandwagons

A bandwagon occurs when a price moves in a particular direction only because it has done so earlier. Agents “jump on the bandwagon” by purchasing the appreciating asset or selling the depreciating one, having no more justification than “everyone else seems to be doing it.”

Bandwagons are another controversial phenomenon in financial markets that can be explained by reference to the Post Keynesian view of forecast determination. They can be seen as a result of availability, anchoring, representativeness, increasing confidence (created by forecast substantiation), and credit/blame issues. As a price movement begins, it is a salient event and therefore more available in memory. Thus, because of the availability heuristic, it plays a disproportionate role in market participants’ forecast—particularly if it was dramatic. Agents may also be expected to anchor to that movement (rather than the level), contributing to the emerging trend as they begin to expect change rather than stability. Representativeness convinces agents that there must be some reason for the movement (if event A is occurring then there must surely be some process B behind it) even if they, individually, do not understand it. Hence, market participants assume that they should be taking the positions implied by the recent price changes. Fourth, as the price continues in the same direction, those forecasting the movement (the numbers of whom may be increasing) will be encouraged by the apparent substantiation of their predictions and thus become increasingly confident and willing to commit funds. Finally, those ignoring the

bandwagon run the risk of being thought foolish, while those jumping on it only to have it go bust can fall back on the excuse, “But everyone else was doing it!” Bandwagons will continue until events (as interpreted by the market participants) indicate otherwise. If those events are sufficiently out of synch with the current run then the confidence of agents may be shattered and a collapse and panic may result.¹⁶

Technical Analysis

Technical analysis amounts to manipulations of past time series of the currency price in order to predict future movements. For example, one may compare a long-term moving average (e.g., the average closing price for the past ten days) with a short-term one (e.g., the average closing price for the past three days). The typical rule is that if the short-term average passes above the long-term, buy the currency in question (and vice versa). Neoclassical economics argues that reliance on such a tool is futile since past price movements are available to everyone. If they contained valuable information then the market would already reflect this. However, Mark Taylor and Helen Allen (1992) found that technical trading rules were very widely used in the foreign currency market. Why would such a simple method be employed by so many professionals if it did not generate profit? From the Post Keynesian perspective the answer is simple: trading rules are profitable, and it is because bandwagon effects make them so. Bear in mind that technical analysis is almost invariably based on the premise that emerging trends will continue (Rosenberg 1996: 324). Any of the various techniques—moving average, momentum, or point and figure, for example—gives a buy signal if recent prices are higher than past ones and a sell if recent prices are

lower than past ones (Schulmeister 1988). If prices continue their climb or fall, then such a rule will result in profit. And if a bandwagon is at work, that is exactly what will happen. Hence, trading rules do not need to rely on self fulfilling prophecy (as some have argued); they simply take advantage of the fact that bandwagons exist. Technical analysis predicts a trend and the bandwagon effect obliges by creating it.

Trading Limits and Cash In

Another factor that can be explained with the above tools is the existence of trading limits in currency rooms. It is a very common practice to restrict the open positions currency dealers may maintain and to set limits on the losses they can incur before they are required to close them out (Weisweiller 1991 and Hudson 1979). This is done because of peoples' psychological attitude toward risk, outlined above. When agents feel they are in a losing situation they will tend to choose more risky options, to "let it ride" and hope that the falling asset price, for example, will turn around at some point. This reluctance to realize a loss can get a trading room into a great deal of trouble, as Barings Bank found out with Nick Leeson in 1995 (Cornford 1996). Hence, limits are placed on dealers so they will not have a choice.¹⁷

The flip side of this is agents' attitude when winning. Under these circumstances, they tend to become risk averse. This is why we observe cash in or profit taking. As a currency appreciates, those holding it become increasingly anxious about realizing what have been up to now only paper profits. Hence, we observe the whipsaw pattern described by Schulmeister (1987, 1988) wherein

even a steadily rising foreign exchange value is continually interrupted by short drops (the moments at which agents cash in).¹⁸

Mental Model

One last factor to be introduced (but not fully developed until chapter five) is the mental model. As suggested in the discussion of framing, the preconceptions of agents are terribly important both in terms of what they consider to be useful inputs into the forecasting and decision making process and in how they believe those inputs will ultimately affect currency prices. These preconceptions, which comprise the mental model, are a social product. Currency market participants are members of a particular subculture. What they believe about how the world works is a direct function of the views of their educators, colleagues, family, friends, kin, et cetera. They seek out the advice of experts, they read professional and scholarly literature, and they consciously and subconsciously mimic the behavior of various role models. We cannot understand the foreign exchange market without understanding what these participants themselves think they are doing.

CONCLUSIONS

This chapter starts the process of developing the tools necessary to understand the foreign currency market. It is argued that portfolio capital flows dominate the market and that within that context market participants' forecasts of future currency price movements are the prime movers

of foreign exchange rates. The model of expectations and decision making (based on psychological research and the economics of Keynes) offered suggests that the market will swing between periods of calm and volatility, that bandwagon effects will be prevalent, and that the foci of expectations formation will emerge as a function of the social context in which the agents interpret their experiences and scholars and professionals engage in research. Once the relationship between exchange rates and the balance of payments is outlined in the next chapter, the following one will use the insights gained to develop a Post Keynesian model of currency price determination.