



BANK FOR INTERNATIONAL SETTLEMENTS

BIS Papers

No 38


FX reserve management: elements of a framework

Claudio Borio, Jannecke Ebbesen, Gabriele Galati
and Alexandra Heath

Monetary and Economic Department

March 2008

JEL classification: F31, G11, G32, E58



This paper was part of the background material provided for an ad hoc meeting on “Challenges in FX reserve management” which took place on 1–2 March 2007. The views expressed are those of the authors and do not necessarily reflect those of the BIS.

Requests for copies of publications, or for additions/changes to the mailing list, should be sent to:

Bank for International Settlements
Press & Communications
CH 4002 Basel, Switzerland

E-mail: publications@bis.org

Fax: +41 61 280 9100 and +41 61 280 8100

© *Bank for International Settlements 2008. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1609-0381 (print)

ISBN 92-9131-755-1 (print)

ISSN 1682 7651 (online)

ISBN 92-9197-755-1 (online)

Foreword

In recent years, issues related to the management of foreign exchange reserves have gained prominence, and reserve management practices have evolved rapidly. Against this background, the Monetary and Economic Department of the Bank for International Settlements organised an ad hoc meeting of senior central bankers from both industrial and emerging market countries responsible for the management of reserves, which took place on 1–2 March 2007. The meeting covered key issues in foreign exchange reserve management, including the broad framework in which it is conducted, asset allocation at the strategic, tactical and portfolio management levels, risk measurement and management, organisational structures and disclosure. This paper was prepared as background and provides a general framework to help organise the various issues that arise in the context of FX reserve management.

Contents

1.	Introduction.....	1
2.	Uses and sources of FX reserve accumulation.....	2
3.	Objectives, risk tolerance and constraints.....	4
3.1	Liquidity.....	4
3.2	Safety and return.....	6
3.3	Constraints.....	8
4.	From first principles to portfolio design.....	9
4.1	Investable assets.....	9
4.2	Horizontal separation.....	9
4.3	Vertical tiering.....	10
4.4	Classifying strategies.....	11
5.	Portfolio implementation and supporting infrastructure.....	12
5.1	Organisation of trading operations.....	12
5.2	Measurement and evaluation of performance.....	12
5.3	Back office.....	13
5.4	Risk management.....	13
6.	Internal governance and public disclosure.....	15
6.1	Internal governance.....	15
6.2	Public disclosure.....	15
	Boxes.....	17

FX reserve management: elements of a framework

Claudio Borio, Jannecke Ebbesen, Gabriele Galati and Alexandra Heath

1. Introduction

In thinking of a framework that might help organise the discussion of FX reserve management arrangements,¹ it is tempting to turn first to the same body of portfolio theory which informs private sector asset management decisions. Indeed, this is what FX reserve managers are increasingly doing in search of clues for how best to discipline and structure their own decisions. This development has gone hand in hand with a change of perspective in reserve management, with the emphasis shifting from liquidity to investment considerations.

That body of work, while inevitably consisting of many strands that are not always fully consistent with each other, does share a number of common building blocks. The first is the set of *investor preferences*, which specify how the investor is willing to trade risk off against return. The second is the *choice of the portfolio* (the set of assets and possibly liabilities) for which risk and return should be traded off. Importantly, this also includes specifying any portfolio items taken as given (“exogenous”) in the exercise, and whose risk profile may influence optimum portfolio decisions (“non-choice items”). The third is the *numeraire* (unit of measurement) in which wealth and hence returns are measured. The fourth is the *investment horizon* over which return and risks are traded off. The final one is the *set of possible investments*, together with their *risk-return characteristics*.

No doubt there are a number of respects in which the portfolio choices faced by a typical asset manager and those faced by an FX reserve manager are qualitatively similar. The determination of the risk-return characteristics of the set of possible investments as well as the choice of the relevant investment horizon are cases in point. In addition, the two types of manager will also face similar choices as regards the implementation of the chosen allocation (eg those relating to how positions are actually taken, monitored and managed).

In the other respects, however, the choices faced by the FX reserve manager appear qualitatively harder, especially once they are embedded in the broader set of trade-offs that a central bank needs to make. First, the FX reserve portfolio, unlike a private sector portfolio, cannot be considered in isolation. It is part of the broader set of assets and liabilities of a country and is held in order to perform a set of economy-wide functions (eg hedging payment needs, insuring against sudden withdrawals of foreign capital). These depend on the broader characteristics of the economy. Second, even abstracting from these factors, FX reserves are just one subset of the public sector balance sheet and, when held by the central bank, of its own balance sheet. The question then arises of the degree to which they should be managed, taking into account other aspects of those balance sheets (eg the proximate financing cost, the impact on central bank capital). Finally, and relatedly, the management of

¹ For the purposes of this exercise, FX reserve management is narrowly defined as the process of managing the stock of FX reserves taking as *given* (exogenous) the net FX positions vis-à-vis the domestic currency and changes in it (eg inflows and outflows of foreign exchange). In other words, the analysis focuses on the allocation of FX reserves across asset classes and instruments, not on decisions concerning their level or on FX intervention. The factors underlying the optimal *level* of FX reserves and those behind their evolution are considered only to the extent that they have an impact on how the given stock is managed.

FX reserves is just one of the many public functions performed by the central bank, raising the question of the relationship between them.

These additional considerations have far-reaching implications. They have a particular bearing on the definition of the risk preferences (or tolerance) of the institution, on that of the broader portfolio in which FX reserves should be regarded as embedded, and on the choice of the numeraire. These choices should ultimately derive from the economic functions performed by the reserves. These considerations imply that traditional portfolio allocation criteria and techniques can only go part of the way towards answering the questions that central banks face. Some of the harder ones are still subject to considerable debate and disagreement.

Given these similarities and differences with typical private sector asset management decisions, this note is organised as follows. Section 2 reviews the economic functions (“uses”) of FX reserves, as well as the reasons behind the accumulation of those reserves (“sources”). With this as a backdrop, Section 3 provides a fuller discussion of the basic principles of portfolio design. It maps the uses of reserves into specific objectives (liquidity, safety and return) and explores the interrelated choices of the specific definition of risk and return, of the numeraire and of the portfolio whose characteristics are considered when deciding on the composition of FX reserves, given the broader constraints under which central banks operate. The articulation of the institution’s risk tolerance deserves particular attention here. Section 4 then proceeds to discuss some of the more technical aspects relating to the design of the portfolio and its implementation. The final section discusses some aspects of internal governance and one aspect of external governance, namely disclosure policies.

2. Uses and sources of FX reserve accumulation

The potential uses of FX reserves can be classified in a number of ways. For present purposes, it might be helpful to consider the following:

1. Intervention in the FX market, with a view to influencing the exchange rate and/or maintaining orderly market conditions.
2. Execution of payments for goods and services for the country, particularly assuming difficulties in obtaining external finance.
3. Granting of emergency liquidity assistance to sectors of the economy, typically the banking sector.
4. Underpinning of investor confidence in the country’s ability to meet its FX commitments, thereby also limiting the probability of financial crises and possibly also reducing the cost of external funding (the “war chest” motive).
5. Execution of payments for the government (possibly within the context of broader debt management operations).
6. Support of domestic monetary policy liquidity management operations (eg through FX swaps, effectively using FX claims as collateral).

The balance between these interrelated objectives will depend on country-specific factors, such as the country’s exchange rate regime, its creditworthiness and its degree of vulnerability to external funding disruptions as well as the range of domestic instruments available for monetary operations. Some of these objectives stress the

precautionary/insurance function of the holding of FX reserves, in which the hedging of nominal and real commitments looms large (eg uses (2) to (4)). Others put more emphasis on normal transaction/working balance uses (eg uses (1),² (5) and (6)). However, with the possible partial exception of the “war chest” motive, they highlight the role of reserves as serving a *liquidity* function.³

To these potential uses it is worth adding an additional one, which arises as a *by-product* of the FX intervention motive:

7. The investment of balances in excess of the foreseeable liquidity/transaction needs.

Table 1
Foreign exchange reserves and measures of adequacy

	Reserves outstanding ¹		Reserves/imports ²		Reserves/broad money ³		Reserves/short-term debt ⁴	
	2000	2006	2000	2006	2000	2006	2000	2006
China	166	1,066	9	16	10	24	8	13
Japan	347	875	11	18	6	15	2	2
Taiwan, China	107	266	9	16	19	34	8	8
Russia	24	295	6	20	44	77	2	5
Korea	96	238	7	9	29	38	2	2
Other Asia ⁵	325	647	6	7	27	30	2	2
Latin America ⁶	136	271	5	7	23	25	1	2
Middle East ⁷	75	178	9	8	25	30	2	2
Central and eastern Europe ⁸	69	181	5	4	39	36	2	1
Industrial economies ⁹	344	334	1	1	3	2	0	0

¹ In billions of US dollars. ² Months of imports. ³ In per cent. ⁴ Ratio; short-term external debt defined as consolidated international claims of all BIS reporting banks on countries outside the reporting area with a maturity up to and including one year plus international debt securities outstanding with a maturity up to one year; based on outstanding year-end positions. For Libya and Saudi Arabia, excludes international securities. ⁵ Hong Kong SAR, India, Indonesia, Malaysia, the Philippines, Singapore and Thailand. ⁶ Argentina, Brazil, Chile, Colombia, Mexico, Peru and Venezuela. ⁷ Egypt, Jordan, Kuwait, Lebanon, Libya, Oman, Qatar, Saudi Arabia and the United Arab Emirates. ⁸ Bulgaria, Croatia, the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Romania, Slovakia and Slovenia. ⁹ Canada, the euro area, Switzerland, the United Kingdom and the United States.

Sources: IMF; Datastream; BIS.

- ² The line is not always easy to draw: FX intervention in the context of an FX and financial crisis has more in common with uses (2) and (3).
- ³ The exception arises to the extent that a country's access to, and terms on, external funding are perceived to depend positively on the stock of net reserves. For instance, since the Asian crisis much attention has been paid to the potential for currency mismatches to exacerbate financial distress, not least by rating agencies. A high stock of net reserves can help to offset currency mismatches elsewhere in the economy, at least if distributional considerations are disregarded. There is some empirical evidence which is not inconsistent with this view. Even so, rating agencies appear to focus primarily on the gross component, stressing the liquidity aspect.

The reason for this use is straightforward. At any given point in time, a country's stock of FX reserves reflects two sources of accumulation. The first is accumulation as a *deliberate* attempt to build the stock up in order to meet potential *future* uses (1) to (6). The second is accumulation as a *by-product* of the implementation of policies aimed at managing the exchange rate and which require adjustments in that stock. In particular, in recent years persistent FX intervention, not least with a view to resisting upward pressure on the exchange rate, has played a key role in many countries. This raises the possibility that the stock of FX reserves may be larger than what might be regarded as desirable for foreseeable future uses (Table 1). Unless there is any FX debt outstanding that can be retired, it will generally not be easy to reduce it without risking "unwinding" the original effect on the exchange rate. This is because a change in gross reserves would also entail a change in net reserves, ie net FX exposure, which would have an impact functionally equivalent to FX intervention.⁴ While, admittedly, the desirable stock in the light of future needs (1) to (6) is very hard to determine with any precision, as the actual stock increases beyond a certain point considerations of the deployment of "excess balances" become more relevant.

3. Objectives, risk tolerance and constraints

The next question is how these uses map into the desirable characteristics of the structure of the portfolio (Diagram 1). Invariably, central banks would note that in their FX reserve management activities they trade off three **objectives**, viz safety ("capital preservation"), liquidity and return. At one level, uses that stress the precautionary and transaction-related motives will generally privilege safety and liquidity; by contrast, the investment of "excess balances" naturally places greater weight on return. Beyond this, though, the usage of the terms can be ambiguous and their precise mapping into the characteristics of the portfolio is not straightforward. Consider, in turn, liquidity and safety/return.

3.1 Liquidity

The term **liquidity** can be thought of as encompassing two related but distinct notions. The first is **market liquidity**. An instrument is said to be liquid if transactions in it can take place rapidly and with little impact on price. This notion in effect covers a form of *transaction costs*, since the costs associated with limited liquidity are only incurred when executing transactions.⁵ The second is **funding liquidity**. This can be defined as the ability to raise funds at short notice, through either the sale of an asset or access to external funding, possibly by using an asset as collateral, in order to meet funding demands.⁶

Uses (1) to (6) are primarily concerned with *funding liquidity*. The main concern is the ability to raise sufficient funds to meet the corresponding demands. Market liquidity does matter, of course, and it interacts with funding liquidity. But, provided transactions can be carried out at

⁴ Of course, the room for manoeuvre will depend on the exact modalities through which the exposures are changed (eg signalling tactics).

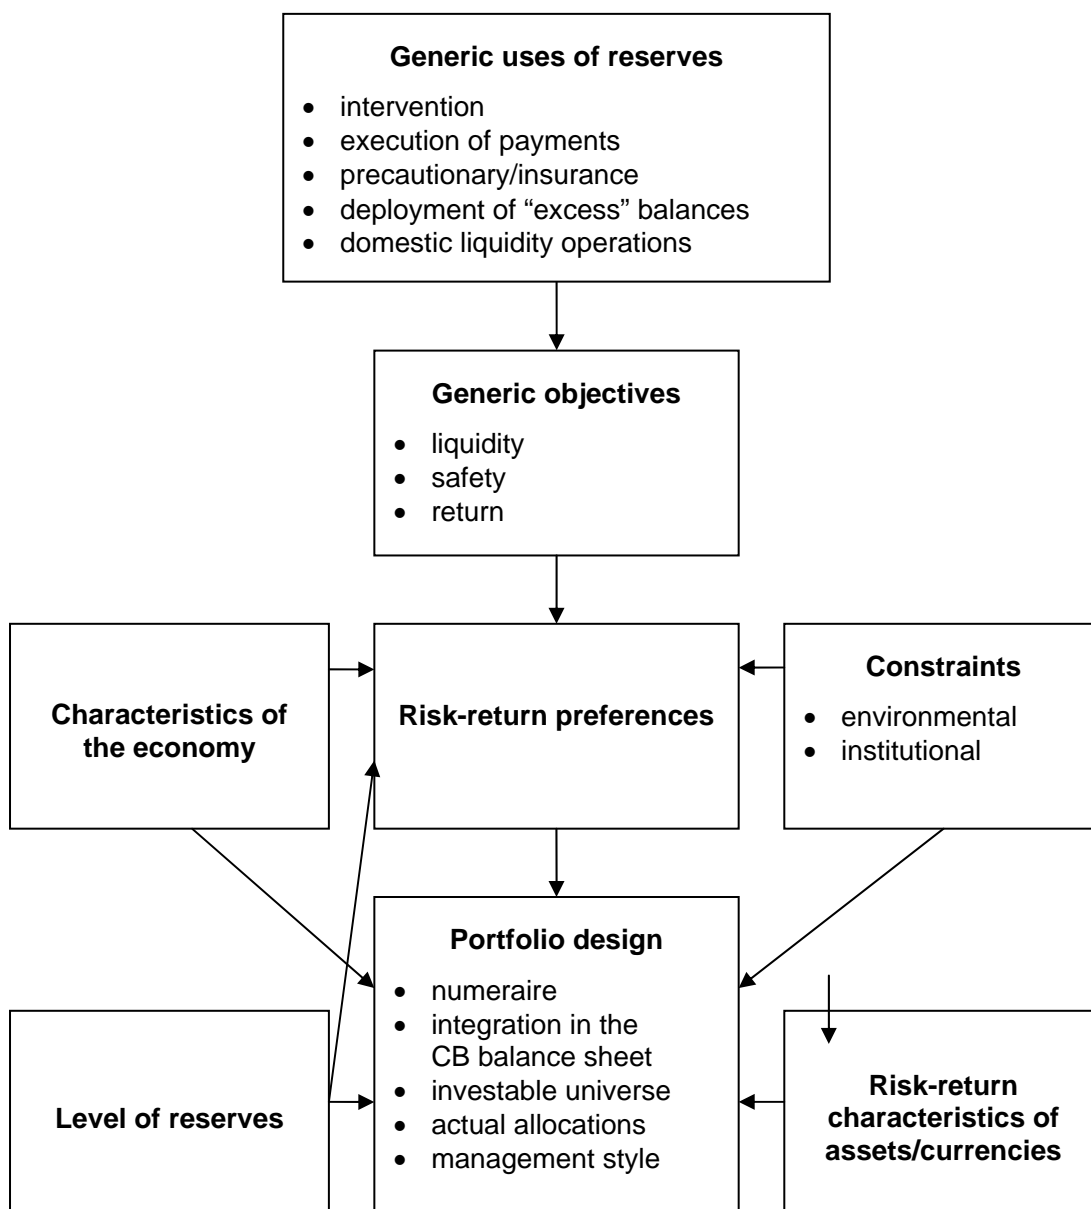
⁵ So defined, market liquidity has several dimensions. Tightness refers to the difference between buy and sell prices, for example the bid-ask spread in a quote-driven market. Depth relates to the size of the transactions that can be absorbed without affecting prices. Immediacy denotes the speed with which orders can be executed, and resiliency the ease with which prices return to "normal" after temporary order imbalances.

⁶ In modern financial markets, funding liquidity is best thought of as including command not only over cash and deposits, but also over other instruments that can be used to meet margin calls and hence effectively settle transactions, most commonly government securities.

short notice, transaction costs are probably best thought of as being incorporated in the net return on the portfolio.

Two implications follow.

Diagram 1
The portfolio design challenge



First, whenever the primary objective of FX reserves is to meet potential liquidity demands, the stock can be built up without incurring any FX risk. The necessary *gross* stock of reserves can be financed in foreign currency (or, equivalently, in domestic currency with a

currency hedge).⁷ Given the characteristics of their liquidity needs and funding costs, some countries have indeed chosen this option.⁸

Second, differential transaction costs are not easily incorporated in traditional portfolio models that focus on risk-return considerations. Bar a clear quantification based on assumptions about the turnover of the portfolio, they tend to be taken into account in the design of optimal asset allocation through qualitative assessments of market indicators, such as bid-ask spreads and volumes, and of asset classes, instruments and currencies. They have, traditionally, loomed large in the context of FX reserves primarily held for intervention or transaction-related purposes (Box 1). Their proper integration, therefore, remains a challenge.

3.2 Safety and return

While liquidity relates primarily to the transaction-related role of FX reserves, safety and return relate to their store of value function. Although concepts such as safety and return may at first sight appear straightforward, their precise definition depends on a number of choices. Admittedly, these choices are sometimes not explicitly articulated, and are hence simply reflected by implication in the structure of the final portfolio. Even so, the increased reliance on quantitative tools in support of asset allocation has tended to make them more explicit.

A first such choice is the unit in which returns and their variability are ultimately measured for allocation purposes, ie the “*numeraire*”. Different choices of numeraire imply different definitions of returns and risk and hence, in general, different optimal allocations.⁹ The choice of numeraire should ultimately derive from an evaluation of the economic function (use) performed by the FX reserves. And it should also take into account the implications of FX reserve management for the other functions performed by the institution, which are “bundled” with it.

To illustrate the nature of the link between the uses of FX reserves and the choice of numeraire, a few examples may suffice (Box 2). Consider first the uses of FX reserves in isolation and treat the FX portfolio as self-standing. If the primary use of the FX reserves is seen as purchasing a bundle of foreign goods and services, then a natural numeraire is the (real) value of that bundle. If the primary purpose is to hedge the value of some of the country’s liabilities, then it could be the currency composition of those liabilities (percentage shares denominated in the various currencies). If the primary purpose is to deploy “excess reserves”, then a unit describing the real purchasing power of domestic wealth might seem appropriate. In addition, if the “location” of the portfolio is not a matter of indifference, other factors may come into play. For example, if the central bank is concerned that losses expressed in domestic currency may weaken its capital base, then the domestic currency would be the relevant numeraire.

⁷ Often, when FX reserves are financed this way, they are referred to as “borrowed reserves”. In a *proximate* sense, though, all FX reserves are “borrowed”, to the extent that they are (generally) backed by forms of domestic debt, as a result of operations aimed at controlling domestic liquidity. In a more *general equilibrium* sense, from the perspective of the *country as a whole*, they need not be “borrowed”. They may reflect the reinvestment abroad of current account surpluses and/or the placement abroad of fiscal surpluses. This would have an effect on the relevant notion of their opportunity cost.

⁸ In principle, it would be possible to rely entirely on credit lines. However, besides typically having a higher cost, they could evaporate when most needed.

⁹ The exception is that of a zero net wealth portfolio in which the objective is full asset-liability matching (risk minimisation). In this case, the choice of numeraire is irrelevant.

What these examples indicate is that the choice of numeraire presupposes an analysis of the ultimate uses and constraints under which decisions are taken. In practice, therefore, it may reflect a compromise between a broad set of factors.

A second, closely related **choice** concerns that **of the portfolio** whose risk-return characteristics are taken into account and, within it, of the items to be treated as beyond the control of the asset manager over the relevant horizon, ie as “non-choice items” or “exogenous”. The gross stock of FX reserves may just be a *subcomponent* of this broader portfolio. Together with the numeraire, the risk/return characteristics of the exogenous portfolio items would help to define the **riskless position** (defined over the relevant horizon (Box 3)) as well as the proximate financing costs of the reserves. The analogy with the management of a pension fund is helpful here. The net wealth of a pension fund would typically be maximised, taking the domestic currency as a numeraire. This would, inter alia, have implications for the hedging of FX risk measured vis-à-vis that currency. In addition, however, the risk profile of the given pension liabilities will play a critical role in determining the riskless position for the portfolio as a whole, as the assets will need to replicate those risk characteristics.

In the case of FX reserve management, in principle various portfolios can be considered. These possibilities can be listed in increasing degrees of comprehensiveness. One is to consider *gross currency reserves in isolation*, regardless of whether they are held against liabilities denominated in FX or not. In this case, the riskless position would be investing the whole of the reserves in an asset yielding a riskless return in the chosen numeraire over the chosen investment horizon. A second possibility is to treat *gross reserves together with any financing denominated in FX*. In this case, the risk profile of the liabilities also sets the risk profile for the gross reserves. Full asset liability management is a strategy consistent with this. A third is to consider also *other investments traditionally regarded as part of a country’s official reserves*, such as those held in the form of gold. The characteristics of gold holdings would then have an influence on the riskless portfolio, which would need to offset the impact on the value of the portfolio of the stochastic properties of the gold holdings. A fourth is to also consider aspects of the *domestic currency balance sheet of the central bank*. This could be taken to determine the cost of funding and/or to help contribute to the impact of the capital and profit and loss of the institution, as partly filtered by accounting conventions. A final possibility could be to also take into account *liabilities/portfolio configurations which are not on the balance sheet of the public entity in charge of the management of the reserves*, but may belong to other entities or the private sector – a kind of enlarged “notional” portfolio. Concerns with the notional hedging of some of the country’s liabilities discussed earlier are a typical example. In this final case, though, precisely because they are not formally part of the physical portfolio over which risk and return are traded off, it is best to think of such considerations as being incorporated via the choice of numeraire.

A third choice concerns the **attitude towards the different risk factors** that may drive the variability of returns and their incorporation into the optimal design of the portfolio – a key aspect of the definition of the all-important **risk preferences (tolerance)** of the institution. In principle, if the agent trading off risk and return was indifferent to the sources of variability, this issue would not arise. Indeed, this is a common assumption in portfolio models.¹⁰ But in the case of FX reserve management, this is typically not so. The implication is that the risk tolerance *needs to be articulated separately* for the various risk factors.

The previous discussion has already considered one source of potential variability, viz market liquidity. Probably an even more important one is **credit risk**. Partly for historical

¹⁰ Box 4 outlines the implications for optimal allocations of another violation of the principle that the return and its variability in terms of the chosen numeraire are sufficient to describe the preferences of the reserve manager, with specific reference to the insurance function of the reserves.

reasons and due to reputation issues, central banks are particularly concerned with the risk of losses arising from the default of obligors. This is true whether it arises on portfolio holdings – the type of risk especially relevant at the level of portfolio design – or in the process of executing the portfolio, reflecting the failure of counterparties in the transactions (see below). At the level of portfolio design, if market and credit risk were fully integrated in risk management systems, one possibility could be to articulate such risk tolerance by setting tolerance levels based on the variability in returns associated with credit risk. A more common procedure, similar to that employed for market liquidity concerns, is to exclude instruments that do not meet a certain minimum requirement (ratings), possibly coupled with floors to the overall credit quality of the portfolio based on those requirements/characteristics of the asset classes.

Thus, the fundamental challenge for central banks managing FX reserves has several dimensions. These include: (i) mapping the uses of the reserves into clear trade-offs between return and risk; (ii) defining the contours of the portfolio relevant for the analysis, including its relationship with other aspects of the institution and the economy's balance sheet; (iii) articulating their risk preferences or risk tolerance; and finally (iv) crystallising those decisions into a portfolio structure and process that are tailored to the institution's needs. Because of differences in country-specific circumstances, therefore, there is no one-size-fits-all.

3.3 Constraints

In reaching any such solution, and as already suggested by elements in the previous discussion, central banks are inevitably subject to a series of institutional and environmental **constraints** that will influence their final choice. At least four such sets of constraints deserve mention.

One set of constraints reflects the ultimate *goals* of the institution. Reserve management operations should not interfere with, and at best should support, those ultimate goals, typically couched in terms of securing monetary stability and contributing to financial stability. For instance, concerns with losses measured in domestic currency should not be allowed to interfere with sound monetary and exchange rate policy. As an example, this could have implications for the choice of target duration and for the desirability of certain variants of balance sheet-wide asset-liability management. Likewise, responsibilities for financial stability may inhibit central banks from investing in certain asset classes (eg if identified as especially risky in their financial stability assessments) or from engaging in inactive asset management along private sector lines (eg implying behaviour which could exacerbate market stress at times of turbulence).

A second set of constraints may result from elements of the *domestic governance environment*. The central bank's relationship with the government plays a key role here. For instance, in a context in which the central bank feels that budgetary independence is important for its operational independence, rules for the distribution of profits to the government, possibly interacting with accounting provisions, can induce the central bank to define comfort levels of risk in relation to its capital. More generally, reputational concerns, paramount in central banks, tend to strongly inhibit risk tolerance. In addition, one may include in this category a variety of constraints reflecting long-standing restrictions embedded in the statutes of the institution; the universe of investable assets can be one such example.

A third set of constraints may reflect the implicit rules of acceptable behaviour by public interest institutions in an *international environment*. In particular, in their FX operations central banks are generally careful not to affect the prices of the instruments in which they operate (other than their own exchange rate). This is not just for reasons applicable to private sector market participants (eg avoiding having markets moving against them when they transact). It also reflects a wish not to disturb prices in those units of account for whose

stability their peers are responsible. This can narrow the range of currencies in which they are comfortable operating and possibly also the scope for active portfolio management.

Finally, constraints may reflect factors *internal to the institution*. These include constraints on the human and technological resources at its disposal, such as on technical know-how and the IT infrastructure. Such constraints operate at all levels of decision-making.

4. From first principles to portfolio design

Moving from first principles to more concrete aspects of portfolio design, three choices merit particular attention. They relate to the set of investable assets and to the degrees of horizontal separation and vertical tiering of the portfolios, respectively. The choice of investable asset delimits the universe of feasible portfolios. The degree of horizontal separation is closely related to the different functions that FX reserves can perform; that of vertical tiering to the desired degree of active management of the reserves, with more active management typically being linked to greater tiering.

4.1 Investable assets

The decision setting the universe of *investable assets* is of fundamental strategic importance: it logically precedes that of specific portfolios and is closely linked to the risk tolerance of the institution. This decision is therefore typically taken at the highest management levels, and is sometimes even beyond the control of the central bank as asset manager, being ultimately the responsibility of the ministry of finance (where the central bank is acting as its agent) or being enshrined in its statutes or other pieces of legislation.

4.2 Horizontal separation

The degree of *horizontal separation* relates to the degree of segmentation in the FX portfolio. One possibility is to manage the whole portfolio as a single unit. Another possibility is to divide it up, based on some form of functional specialisation. The most common distinction here is between “liquidity” and “investment” tranches. In practice, the line between these two tranches can be drawn in different places. In some cases, it is drawn between working balances and the rest; in others, between the more immediate liquidity needs and the rest. Another possible distinction is between such tranches and a “liability” tranche, designed to hedge a specific set of liabilities.

What these demarcations have in common is that they segment the portfolio according to different *uses* of the reserves. By so doing, they can yield gains in terms of clarity of purpose and focus.¹¹ In addition, given the different implications of different uses of the reserves for the objectives and basic inputs into the construction of the portfolio, such as the choice of numeraire, segmentation can in some respects simplify the actual construction of the targeted portfolio. On the other side of the ledger, though, such segmentation forgoes the possible synergies of a more integrated approach. As the size of the portfolio increases, not least as a by-product of intervention policies, the incentive to segment the portfolio along functional lines may also increase.

¹¹ They may, however, reflect broader institutional factors, such as the origins of the various portfolios.

4.3 Vertical tiering

The degree of **vertical tiering** relates to the number of layers in the design and implementation of the portfolio. Here the main distinction is between a three-tier and a two-tier structure.

As the degree of discipline and rigour in reserve management has increased, the design of portfolios will almost invariably have a critical top level (the “**strategic asset allocation**” (SAA) level). It is at this level that the key characteristics of the portfolio are defined. This decision will take the form of a specific point target allocation to asset classes (FX, types of fixed income, possibly other asset classes) together with reproducible indices to measure performance (jointly referred to as the “**strategic benchmark**”).¹² In addition, these point targets will normally be complemented with *tolerance ranges*, setting permissible deviations from the corresponding target allocation for the next management tier and hence implicitly also the permissible scope for active strategies. This level is generally thought to determine over 90% of the return of the portfolio. It may be dispensed with in the case of some small liquidity tranches. While investment guidelines may still constrain allocations in such tranches, they may be simpler and less precise and, above all, not require the explicit measurement of performance.

All structures have in common a bottom level at which positions are actually taken in the market (the “**portfolio management**” (PM) level). Because portfolio managers typically have some freedom in the positions they take, to a certain degree they are in fact not just implementing the portfolio but actually contributing to its design, albeit in a *bottom-up* fashion and often within strict guidelines.

Finally, in a number of cases an intermediate level exists (the “**tactical asset allocation**” (TAA) level). The TAA will generally be specified through point target allocations that deviate from the strategic benchmark (the “**tactical benchmark**”). In addition, it will be complemented by permissible tolerance ranges around them, constraining the portfolio management position-taking.

The objective of such an intermediate level is twofold. It is to centralise portfolio decisions at a higher level of management, where the “bigger picture” expertise lies, allowing a more consistent “view” to be formed of how to obtain higher returns in relation to the strategic benchmark. And, in the process, it is to enable a higher degree of active position-taking, defined by the deviations from the strategic benchmark, within a more disciplined framework. Not surprisingly, this intermediate level has become increasingly popular as return considerations have grown in importance.

In principle, the main distinction between the SAA level and the others is that the SAA is intended to set the longer-term allocation guidelines for the portfolio, *without* taking a view about the shorter-term movements in markets and seeking to profit from them. Other levels are designed to exploit such shorter-term movements. Sometimes, this difference is summarised by stating that the investment horizon of the SAA is longer than that of the other levels. In practice, however, it may be hard to draw a sharp distinction between taking and

¹² At a strategic level, reliance on derivatives can allow the decoupling of positions in the underlying cash markets/asset classes from the net exposures in the portfolio. In particular, an *FX overlay* permits the definition and implementation of a strategic currency benchmark which is different from the summation of the positions through the underlying cash instruments (eg bonds and deposits). At the tactical and, possibly, portfolio management level, FX overlays can also be used as part of active management, with a view to beating the strategic benchmark.

not taking a view, depending on the specific nature of the inputs used and the precise definition of “horizon” (Box 3).¹³

4.4 Classifying strategies

For the portfolio as a whole, it is then possible to describe portfolio strategies in two dimensions (Table 2). The first is the degree of *tightness of the deviations* from the strategic benchmark. The second is the degree of *stability in the strategic benchmark itself*. At one end of the spectrum, stable benchmarks with small deviations imply very stable allocations (in value terms) and hence frequent *rebalancing* of the allocations towards the benchmark (“beta-driven” style). Indexing to the benchmark is the extreme case. At the other extreme, time-varying benchmarks with wide deviations imply a very active management style, with a view to obtaining additional returns (“alpha-driven” style).

Clearly, too, a given degree of active management may be achieved in two ways. One is to adjust the benchmark but have smaller tolerance ranges around it (a more “top-down” approach); the other is to take most of the adjustment in the deviations (a more “bottom-up” approach). The intermediate tactical level between the strategic and portfolio management levels adds one degree of freedom and may be particularly useful in those cases in which the stability of the strategic benchmark is regarded as important.

Table 2
Types of management styles

		Tolerance ranges ¹	
		Narrow	Wide
Stability of the benchmark ²	High	Top-down Passive ³	Bottom-up Active
	Low	Top-down Active	Bottom-up Active

¹ Tolerance ranges around the specified benchmark. ² This refers to the strategic benchmark, but is in principle also applicable to the tactical benchmark when considering the relationship between that benchmark and portfolio management. ³ If the tolerance range is very narrow, this effectively amounts to indexing to the benchmark.

The implications of a buy and hold strategy will partly depend on the asset classes considered. On the one hand, the strategy, itself characterised as very passive, could be regarded as reflecting the absence of a stable target SAA and as implying no rebalancing. In this case, portfolio weights would drift passively, reflecting market movements and the criteria used for the marginal allocations. On the other hand, if the flows are sufficiently large

¹³ Sometimes the distinction between the SAA and other levels is described in terms of the source of returns. In this view, the performance of the strategic benchmark reflects exposure to systematic market factors (or “beta”), while that of the tactical benchmark and any active portfolio management away from those benchmarks captures gains and losses that are uncorrelated with those factors (“alpha”). More precisely, one can think of the strategic benchmark as defining the systematic risk factors or “market portfolio” with respect to which the performance of the tactical benchmark and the actual portfolios is measured. However, when the strategic benchmark is not based on unconditional estimates of the statistical parameters (eg expected returns, variances and correlations), and is revised in line with expectations of shorter market trends, the distinction is less clear-cut.

and the desired deviations from benchmarks small, one can implement a target fixed income allocation, such as a target duration, via adjustments in the flows, including those arising from maturing assets. This would be harder for currencies, though, given their typical volatility.

5. Portfolio implementation and supporting infrastructure

Implementing the designed portfolio involves several further stages. These include: (i) the organisation of trading operations (the PM level); (ii) measurement, monitoring and evaluation of performance; (iii) the actual processing and settlement of trades (back office functions); and (iv) the management of the risks incurred in implementation (market risk, credit risk, operational and legal risk).

5.1 Organisation of trading operations

The organisation of trading operations involves a number of choices. One concerns the specific strategies allowed. A second concerns how the incentive structures of traders are determined through the types of mandates provided, the response to underperformance, if any, and other aspects of rewards and penalties. A third concerns the choice of whether to manage the portfolio in-house or give at least some segments to external managers. As the universe of assets in which central banks invest has broadened and return considerations have become more prominent, the recourse to external managers has increased. Key questions here relate to the criteria used for deciding to outsource (eg gain access to asset classes in which know-how is missing, benchmark internal performance, help develop expertise), for choosing among competing offers (eg costs and performance, strength of internal processes, ability to provide training) and for evaluating their performance ex post.

5.2 Measurement and evaluation of performance

The measurement and evaluation of performance consists of a number of steps: the valuation of the portfolio and calculation of the returns, the attribution of those returns to different decisions and the evaluation proper.

Measurement relies on proper **valuation** procedures. Regardless of whether portfolios are managed in-house or externally, processes need to be in place to value portfolios. As FX reserve management has become more sophisticated, marking to market for management, as opposed to external reporting, purposes, has naturally become the rule. This has increased the demands on the technical infrastructure and on the processes aimed at reconciling these valuations with those for public accounting purposes. Depending on the nature of the instruments and the investment process, valuation may be outsourced, at least in part (eg to custodians or external managers). Moreover, as the instruments traded become increasingly complex, fair value measures may require the use of “marking to model”, possibly raising questions about the accuracy of the underlying valuations. The accuracy of valuation also depends on the frequency of measurement. In general, the higher the frequency of valuation, the more accurate measurement can be, as it allows a better separation of returns from any reserve inflows/outflows. A higher frequency can also support the implementation of checks on position-taking, by resulting in more accurate and timely monitoring.

Performance **attribution** can be thought of as having two dimensions. On the one hand, attribution simply breaks down performance “vertically”, into the various decision levels (eg actual performance versus the strategic benchmark; tactical versus strategic; portfolio management versus tactical or strategic, depending on whether a tactical level exists). This

is a relatively simple exercise, consisting of calculating the difference between the returns achieved in the actual portfolio and those of the benchmarks. On the other hand, and more demanding, it can break down performance “horizontally”, into the individual (risk) factors that drive return. In this regard, over time systems have become increasingly sophisticated, being capable of attributing performance to increasingly “granular” factors, down to security selection. How far a central bank decides to invest in this area will depend on several factors, including the extent to which active management is pursued as well as the availability of technological and human capital.

The **evaluation** proper calls for several choices. One is the specific metric used. Evaluation may be based simply on the return, but risk-adjusted measures have become increasingly common (eg Sharpe ratios). Others relate to the frequency of the evaluations and the period over which performance is actually evaluated. These trade off accuracy in measurement, limiting “noise” as the period lengthens, against the risk of accepting any underperformance for too long.

5.3 Back office

The processing and settlement of the trades are the final stage of the trading process. This is the main function of the **back office**. In addition, the back office may perform a variety of other activities. These may include accounting, confirmation of trades, data management and dissemination, collateral management and IT system support. A key issue here is the degree of automation. Not only does automation reduce the marginal costs of trading, improve its efficiency and better support its growing volume. Increasingly, it is also seen as critical in limiting the exposure of central banks to financial and reputation losses arising from (unintentional or intentional) human error. As a result, it has received great attention in the context of operational risk mitigation, particularly through a higher reliance on straight through processing (STP; see below). A second issue, as with portfolio management functions, is whether there is a cost-benefit case for outsourcing these activities. A third issue, shared with other areas of reserve management, is how to deal with new business requirements, such as developing knowledge of new instruments and of markets in-house.

5.4 Risk management

The management of the risks incurred in the implementation of the portfolio takes place at various stages.

The assessment of (ex ante) **market risk** is an integral part of the formulation of the SAA and it clearly plays an important role in that of the TAA, if any, and of the PM positions. When setting the SAA, it is the overall (*absolute*) market risk of the portfolio that is relevant; at the other levels, it is the market risk relative to the corresponding benchmark that matters. Mandates at the various levels set tolerance ranges which implicitly or explicitly trade off the potential for return against such market risk.

A number of issues are particularly relevant. One is *measurement*. A range of tools can be employed to measure such risk, depending on the specific portfolio involved and risk technology available. These range from variants of duration/gap analysis (for interest rate risk) to variants of VaR, stress tests/scenarios and even full ALM simulation techniques. A key question here is what the relative merits of the various approaches are. In this context, backtesting acquires particular significance, as a means to evaluate the performance of the models used and of their inputs. Relevant questions in backtesting include the level of aggregation over subportfolios, the frequency of the exercise and the profit and loss measure used, which will in turn depend on the frequency of valuation. A second issue concerns the *management* of risk proper. This relates to the frequency of monitoring and to the tools used to set various risk limits. It is sometimes possible to measure risk with one tool (eg VaR) but

to set limits with another (eg duration), depending on factors such as simplicity and enforceability.

The management of **credit risk** is a critical and challenging area: despite the particular sensitivity of central banks to credit risk, techniques are not as refined as those for market risk and several open issues remain. Credit risk can be thought of as reflecting two sources. The first is the credit risk associated with the failure of the issuer in whose instruments the portfolio is held ("**portfolio credit risk**"). These can be deposits, primary securities and, increasingly, derivatives – items that also vary in terms of collateralisation. The second is the credit risk incurred in the process of executing the transactions required to implement the portfolio, vis-à-vis the institution which is facilitating them ("**transaction credit risk**"). Settlement risk, reflecting either no simultaneous settlement of the two legs of the transactions (no DVP or PVP) or lags between trade and settlement dates, is the cleanest example. Unfortunately, the often used distinction between "investment" and "counterparty" credit risk only partly captures the dividing line between these two sources, as the terms are used differently across central banks, complicating comparisons.

Regardless of the source, a number of issues arise in the context of credit risk management (Box 5). One is the basis for the *credit analysis* of individual issuers and counterparties in transactions. This can range from exclusive reliance on rating agencies to a combination of such assessments and own analysis, based on balance sheet and/or market information (eg from prices). A second is the *measurement* of the corresponding exposures. Relevant questions here include, in particular, the valuation method, the range of exposures to single entities that are included (eg the treatment of settlement exposures) and the methods of aggregation across different entities (eg some form of addition or methods that take into account correlations). A third is how *limits on exposures* are set (coverage and criteria), revised and enforced, in turn depending on the frequency of monitoring. A fifth covers reliance on *other risk mitigation* devices, not least collateral policies. A final one concerns the *relationship between the management of counterparty credit risk in reserve management and domestic monetary policy operations*.

Operational risk is perceived as an area of key and increasing importance. Although the Basel Committee has defined operational risk as "the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events", risk of financial losses is closely linked to reputation loss for central banks. While operational risk has always been managed to some extent, there has been a clear trend towards a more comprehensive treatment, analogous to the way in which financial risks are managed.

As with any other type of risk, a first issue concerns *measurement*. Lack of data makes it particularly difficult to derive effective quantitative measures of vulnerability to operational risks. As a result, it is more common to rely on simple indicators. These include the size and complexity of the portfolio, the degree of independence given to individual portfolio managers, the concentration of knowledge, experience and skills in specific individuals, and the degree of STP. A second issue concerns the *management* of the risk proper. For example, **legal risk**, which arises from the possibility that financial losses will ensue from not conforming to national and international laws, may be mitigated by using standard contracts and ensuring that internal processes and systems are well documented. Examples of activities that would help to mitigate the operational risks arising from process failure include reorganising the different stages of the trading process to ensure that the reporting lines for functional areas with potential conflicts of interest are separated (see below), and increasing the degree of automation. Operational risk arising from external events can be mitigated by putting in place detailed business continuity plans, including by setting up remote system recovery sites.

6. Internal governance and public disclosure

6.1 Internal governance

Over time, internal governance issues have risen within the agenda of central banks. Many factors have combined to strengthen this trend. Some of these factors reflect changes in the broader institutional environment in which central banks operate. The shift towards greater independence, and hence the need for greater accountability, is one such example. Strengthened *internal* governance should be seen as the natural counterpart to strengthened *external* governance. Other factors are more specific to developments within the reserve management world. In particular, the trend towards greater emphasis on return has called for a more structured and top-down approach to reserve management. These considerations have simply added to the traditional incentives of central banks to limit reputational risk generally. The concern here is that a loss of reputation through bad governance could undermine their ability to perform their primary tasks of ensuring monetary stability and contributing to financial stability.

Two areas of internal governance merit particular attention. They may be referred to as “vertical” and “horizontal” governance, respectively.

One of the core principles of good governance is that decisions be taken at the right level of the organisation (“**vertical governance**”). In particular, it is recognised as important that the executive level of the organisation be responsible for strategic decision-making: for example, in making decisions about the risk tolerance of the organisation. The concern is that otherwise such key decisions could be taken without proper guidance from the top.

Another core principle of good governance is to organise business areas and their reporting lines in a way that minimises the potential for conflicting incentives, particularly when these conflicts could result in behaviour that is reputation-damaging, if not fraudulent (“**horizontal governance**”). The key issue here is to avoid or limit conflicts of interest in the exercise of the various functions, largely through the horizontal separation of activities. For example, at the portfolio management level, any incentives traders may have to achieve higher returns by misrepresenting trading activities can be reduced by separating the trading function from the performance measurement, risk control and compliance functions. At higher levels, the main concern may be that access to private information concerning monetary policy decisions, of the central bank that is managing the reserves or of others, should not be allowed to inform portfolio choices.¹⁴

As in all areas in which conflicts of interest arise, the key issue is how to balance reputational risks with the loss of potential synergies (economies of scope) in the joint performance of the various functions. This raises issues of optimal organisational structure that inevitably need to take into account a series of factors, including the size and availability of human capital in the institution. For example, it may be optimal from a governance point of view to ensure that the reporting lines of the front and middle offices remain separate up to the executive level, but the specialisation and experience levels of management below this level may make it more sensible for reporting lines to intersect at a lower level in the central bank.

6.2 Public disclosure

Disclosure plays a dual role. On the one hand, it is an element of external governance. FX reserves are part of the assets of a country, and the central bank manages them in the public

¹⁴ This means that horizontal separation may not be enough: in some cases, it is also a question of not having senior people with access to such information be involved in active management (eg at the tactical level).

interest. As such, the central bank may be called upon to disclose how it fulfils this role as part of governance arrangements. On the other hand, the extent and type of disclosure can have an impact on how effectively the central bank discharges its responsibilities for monetary and financial stability.

From both perspectives, developments have encouraged a trend towards greater disclosure. Tighter external governance naturally means greater disclosure. In addition, on balance, central banks have become more favourably disposed to the view that greater disclosure can help in the performance of their primary responsibilities. For example, in the specific FX reserve area, following the Asian crisis, the international community has agreed on a strengthened disclosure standard, encapsulated in the IMF's Special Data Dissemination Standard (SDDS). Such a standard, though, relates only to the level of reserves and some aspects of their composition. It is silent about most of the main aspects of FX reserve management, notably the allocation across currencies and fixed income assets. As a result, disclosure practices are left to the discretion of national authorities, reflecting a mixture of economic, institutional and political factors.¹⁵

In the past, much of that discussion about disclosure was concerned with the overall *level* of stock of FX reserves and of the transactions that affect changes in that stock (intervention), not its *composition*. Often, the concern was that limited disclosure could increase the risk of overestimating the stock of reserves of a country, thereby reducing the discipline that financial markets could exercise on it. More recently, though, as reserves have grown, attention has switched to their composition – a type of information which, traditionally, central banks have been less inclined to provide. The argument put forward has been that greater disclosure, if properly structured, can actually help to stabilise markets.

External governance arrangements aside, and at the risk of some oversimplification, the optimum degree of disclosure has been seen as trading off the efficiency-enhancing effect of the provision of additional information to financial markets, on the one hand, against the loss of tactical room for manoeuvre for the central bank, on the other. Differences of view, therefore, hinge on differences in perspective on how well markets function and on how much disclosure actually constrains the authorities' ability to pursue their objectives. Factors that may constrain such room for manoeuvre include the size of the reserves and the nature of the exchange rate regime. If the FX reserves are very large, then the assumption that the central bank is an atomistic player in the international currencies and in the instruments held in its portfolio is less likely to hold. If the regime is one in which the exchange rate is heavily managed, then tactical considerations can play a bigger role. For instance, information about changes in the currency composition of the reserves may be seen as foreshadowing changes in exchange rate regimes. This is so to the extent that the composition is related to the reference currency (eg a shift from a bilateral to a basket reference currency). Similarly, this type of information can also help disentangle ex post changes in the level of reserves that reflect intervention from those that reflect valuation effects.

¹⁵ The IMF has also produced a set of guidelines for foreign exchange reserve management, which include transparency criteria. These, however, are quite general and allow ample room for discretion.

Boxes

Box 1

Incorporating market liquidity considerations into the FX portfolio

In contrast to safety and return, which relate to the store of value function of the reserves, market liquidity relates to their role as a means of payment. Minimising transaction costs associated with differential market liquidity conditions (across both instruments and time) involves limiting the costs of (a) liquidating the underlying instrument and, when necessary, (b) executing the trade of one currency for another.¹ Step (b) is always involved in the case of FX intervention when the foreign currency is traded for domestic currency. It would also be involved whenever the currency in which the reserves are originally held is not the one needed for the transaction, even if the domestic currency is not part of the exchange at all. One such example would be liquidating an asset held, say, in dollars to extend a loan or to pay for a service denominated in another currency, such as the yen or the euro.

Historically, the dollar has been the preferred currency on liquidity grounds. It tends to have some of the most liquid markets in the underlying instruments and to trade at the narrowest spreads in the FX market, as the vehicle currency par excellence. In fact, the statistics from the BIS Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity for 2007 indicate that the dollar is one side of 86.3% of all FX transactions, followed by the euro and the yen, which enter, respectively, in 37.0% and 16.5% of trades. Because of this advantage, it may often pay to hold the dollar or another very liquid currency, even if the ultimate liquidity needs are denominated in another.

At the same time, underlying markets in other currencies, such as the euro, have gained ground in recent years in terms of liquidity. And it is not always true that the most liquid FX segment in which the domestic currency is traded is the US dollar. For example, in some countries which have as their reference foreign currency the euro, the euro segment may be the most liquid. Some markets for European currencies are one such example. More generally, whenever the reference currency and the most liquid FX segment coincide, it would be empirically hard to distinguish between the liquidity and safety (value preservation) motives for holding reserves in that currency. The reason is that the reference currency is also the one vis-à-vis which the volatility is, by definition, lowest when measured in domestic currency (Box 2).

Liquidity concerns would generally limit the number of currencies held in a portfolio as reserves, especially if considered alongside declining diversification benefits at the margin. This helps to explain the dominant role of a few international reference currencies. The more political concerns with the possibility of moving markets in a unit of account for which a peer institution is responsible add to this constraint. At the same time, with the development of more liquid derivatives markets, sometimes a central bank could have a currency exposure to one currency without actually holding any underlying assets. This could help to reduce the price impact of the transactions.

Liquidity considerations would generally be reflected in a portfolio through specific constraints. These would set ceilings to assets thought to be relatively illiquid and/or floors to currencies or asset classes regarded as particularly liquid. An alternative would be to rely on models that explicitly trade off return against liquidation costs. This would emphasise transaction cost considerations over variability of holding period returns.

¹ For present purposes, transaction costs would also involve any delays or, in the limit, the impossibility of executing transactions.

Box 2

The choice of numeraire

The numeraire is defined as the unit for the measurement of value in a given portfolio. It therefore plays a critical role in any formal asset allocation problem. It helps to define actual and expected returns on a portfolio, and it also helps to establish the riskiness of those returns. In particular, given risk aversion, optimal portfolio allocation penalises variability with respect to the numeraire.

Because of its function, arguments for or against the use of particular units of measurement as an input into the portfolio optimisation process have to be derived from first principles, based on the functions of the reserves. These, in turn, can be used as a proxy to infer the preferences (utility function) of the reserve manager. The numeraire can be defined in nominal or real terms. It can also be set in terms of individual currencies or baskets. What follows considers some possibilities and their implications.

If the intended use of the reserves is to ensure access to imports under stress conditions, a natural candidate for the numeraire is the corresponding *basket of imports*, measured in real terms. The clearest analogy here is with a consumer, who derives utility from consuming a basket of goods and services and derives disutility from variability in the purchasing power defined in terms of that basket. The shares of the currencies in which imports are denominated would be the relevant concept for the (short) horizons over which prices are fixed (predetermined) in those currencies, as it would capture the cost of the basket. Over longer horizons, though, the correlations between exchange rates and the underlying prices would need to be taken into account.¹

If the intended use of the reserves is to hedge the value of a (subset) of a country's liabilities (debt) out of, say, concerns with limiting the risk of a financial crisis, then a natural candidate for numeraire is the *currency composition of those liabilities*. This would mean that, in response to changes in exchange rates, the value of the numeraire would change in proportion to that of the corresponding liabilities, thereby stabilising the value of the reserves in relation to them.

If the intended use is to intervene so as to influence the value of the currency, the numeraire currency is less obvious. A possible criterion would be to use the *currency regarded as most effective for intervention purposes*. If this was thought to depend on the currency or currencies with reference to which intervention was pursued, then these reference currencies could be used. In general, however, the currency of intervention would be chosen based on the transaction costs involved, which would depend crucially on the liquidity of the FX market for that currency and on that of the underlying investments (Box 1).

Finally, if the intended use for the FX reserves is the deployment of funds in excess of foreseeable liquidity needs, then a possible candidate for numeraire would be the basket of goods and services, presumably measured in real terms, which most closely represented *the deflator for the country's "wealth"*. Inflation aside, this criterion would attach a high weight to the domestic currency.²

The domestic currency is also a natural choice for numeraire if the central bank is concerned about the impact of fluctuations in the value of the reserves on its profitability and capital (ie not so much based on the functions of the reserves as on the consequences of holding them for other functions performed by the central bank). In this case, stabilising the return on reserves would imply stabilising it also in terms of domestic profits and capital.³

Whenever the domestic currency is chosen as the numeraire, the nature of the exchange rate regime has a direct impact on the currency allocation. In the case of bilateral pegs, for instance, using the domestic currency would be equivalent to choosing the reference foreign currency as numeraire. Ceteris paribus, this would tilt the composition of the reserves towards that currency. Likewise, in the case of a basket peg, the corresponding basket would play the same role. More generally, FX reserves would be tilted towards whatever composition of foreign currencies was more stable relative to the domestic unit of account.⁴

¹ For example, in the extreme case in which purchasing power parity holds over the relevant horizon (without any uncertainty), the weights would be undefined/immaterial, as all goods would be perfect substitutes. The result is analogous to what happens when the numeraire is defined in nominal terms and exchange rates are fixed, in which case all currencies would be perfect substitutes (liquidity and credit risk considerations aside). ² Of course, in formal terms, the maximisation exercise would need to impose the constraint of taking the overall net position vis-à-vis the domestic currency as given. Otherwise, the risk-minimising position would tend to assign a high weight to the domestic currency (ie imply a reduction in the overall net reserves). ³ This assumes mark to market accounting. See below for a further treatment. ⁴ If the strategic asset allocation exercise assumed that expected returns across currencies are equal – a common, neutral assumption in long-term, unconditional analysis and equivalent to imposing uncovered interest rate parity – then the (unconstrained) currency composition of the portfolio would be identical to such a risk minimisation basket. In the extreme case of pegging, it would thus be equal to the reference currency (basket).

Box 3

Horizons in portfolio management

The notion of “investment” horizon plays a key role in the theory of portfolio management. This notion is important in defining the “riskless” returns and the risk tolerance of the investor. In addition, it can be related to the degree of active management. At the same time, the term “investment horizon” is quite ambiguous; as normally used, it can in fact conflate a number of different concepts. And the relationship with risk tolerance and the degree of active management is not immediate.

A number of different notions of horizon can be distinguished.

First, the **holding period** or **planning horizon** is the period over which return and risk are actually defined. It thereby also defines the riskless asset which dominates pure “cash”, as the asset that yields a *deterministic* return in the relevant numeraire *over that interval*. Thus, if the holding period is, say, one year, then inputs concerning returns and measures of risk/variability would be calculated over that horizon.¹

Second, whenever the asset manager is *not* indifferent to fluctuations over the relevant planning horizon, it is necessary to consider also an **intermediate risk horizon**. This is the shortest interval over which the risk measure is defined. For example, while the planning horizon may be *a year*, the exercise may set an additional constraint requiring that losses not be expected to exceed a given level with a given probability *in any given month*.

Third, the **revision interval** is the interval between the updating of the assumptions relating to the *statistical properties of the returns* expected to prevail over the relevant planning horizon. Ordinarily, this would be expected to be longer than the planning horizon. It might be odd to revise views about returns, variances and covariances before observing at least one full realisation of them. At the same time, revisions could also occur in response to market developments that challenge those assumptions even before a whole cycle of observations has taken place (“state contingent” interval).

Fourth, the (framework) **review interval** is the period between revisions of the framework or elements thereof. Again, here one would expect a mixture of fixed (normal) intervals combined with flexible intervals in the light of relevant market developments and other factors.

The *risk tolerance* of the investor can be regarded as being measured by, say, the size of the acceptable loss over a given time interval. *Ceteris paribus*, the longer that time interval, the higher the risk tolerance. If the investor cared only about returns over the planning horizon (holding period returns), then that horizon would be sufficient to define its risk tolerance. When the investor also cares about fluctuations within that interval, then in effect its tolerance for risk is lower, as its effective horizon is shorter.

The *degree of active management* depends more directly on the frequency of asset reallocations in the light of revisions in the assumptions about the statistical properties of the returns over the relevant planning horizon. Long planning horizons with short revision intervals will reflect more active strategies than short planning horizons with long intervals.

¹ It is not uncommon, however, for exercises to take as the riskless return (risk-free rate) a short rate (eg a three-month rate) even if the planning horizon is longer (eg one year).

Box 4

Taking the insurance role of FX reserves one step further

The typical portfolio analysis implicitly assumes that, for a given variability of the return in terms of the numeraire, this return is sufficient to describe the “utility” that the investor derives from the portfolio. In other words, the FX manager would be indifferent between two portfolios yielding the same return/variability configuration. As already discussed in the text, liquidity and credit risk considerations are examples that violate this principle, but they are not the only ones.

The FX manager may also care about the *background circumstances* against which a given value of the portfolio materialises, because those circumstances affect his/her utility, more broadly defined. The role of FX reserves as “war chest” against potential financial crises is an important such example. The FX reserves can be more effective in providing such insurance if their value is highest precisely when the probability of a crisis given background circumstances is highest.¹ This will maximise the capacity of the stock of reserves to reduce the cost of any crisis that may materialise and may also maximise its contribution to limiting the probability of the crisis occurring in the first place. This is what insurance is all about.

This perspective has implications for the optimum allocation across asset classes and currencies. The basic principle would be to invest in assets/currencies that are likely to rise in value as a crisis looms or materialises, thereby maximising the reserve manager’s ammunition to fight it.

To some extent, the way reserves are typically invested already performs this function. This is so whenever they are held in assets/currencies that act as “safe havens” at times of stress, thereby tending to appreciate in value. Treasuries and gold are such examples.

Going beyond this requires having some idea of what the possible causes of the crisis might be. Consider a few examples. For a country that is highly dependent on export revenues from particular commodities, a possibility might be to use part of the portfolio to short those (or related) commodities. Alternatively, if the origin of the crisis is assumed to be a widespread reversal in global investors’ risk appetite, then holding short positions in emerging market assets might be another possibility (eg in the CDS market). In all of these cases, as long as the corresponding markets are, and remain, sufficiently liquid and the reserve manager acts as a price taker, these strategies would be *technically* sound.

Of course, this form of insurance, like any insurance, could well come at a cost. If the relevant asset had *generalised* safe haven properties, it would be priced accordingly: with risk-averse investors, its yield would tend to be lower *on average* (and especially in good times) in order to compensate for its high value in bad times. But if a country is vulnerable to *idiosyncratic* risk factors, this would not generally be the case. A commodity like oil or copper, for instance, is an export for some countries but an input of production for others.

Arguably, the main problems with an approach of this type are of a political economy nature. For a country which produces a particular commodity, it would be hard to “sell” a strategy that could be perceived as betting against its own producers. And the feasible amounts may not be large enough to provide benefits that outweigh such political costs. Even so, a further exploration of this approach may be useful. At a minimum, it suggests that surveying the universe of investable assets and currencies in search of those that might have insurance properties tailored to specific countries could be a complementary approach to the more traditional one.

¹ More technically, this means when the probability of a crisis *conditional* on the background conditions, but not on the value of the FX reserves, is higher.

Box 5

Credit risk management

Regardless of the basis of the credit analysis performed by the central bank to assess the riskiness of its obligors (ie those entities against which credit risk is incurred), a number of choices need to be made. They concern: (i) the degree of integration of the management of various types of credit risk in FX; (ii) the measurement of exposures to individual obligors; (iii) the aggregation of exposures across obligors; (iv) the nature and management of the limit systems in place; (v) other risk mitigation devices such as collateral; and (vi) the degree of integration with other operations carried out by the central bank, notably those involved in domestic monetary policy implementation.

Degree of integration within FX reserve management

In principle, there are various ways of addressing credit risk exposures arising from “portfolio” and “transaction” credit risk. At one extreme, one can imagine a fully integrated system, in which portfolio and transaction risk are treated as dimensions of a global system, used to measure the exposures and to set limits across them. At the other extreme, one could imagine segmentation by instrument and type. In general, the main dividing lines in a less than fully integrated system could be drawn in various places. This would partly depend on the degree of overlap between obligors under the various categories and the nature of the measurement and management systems in place.

One possibility could be to draw the line between settlement and other risks. The risks associated with the lack of synchronisation of the two legs of the transactions can be large but are also much shorter-lived compared with those arising from the holding of instruments in the portfolio. They can also be addressed through specific mechanisms, such as delivery versus payment (securities) or payment versus payment (FX).

Another possibility is to draw the line between what is sometimes called “investment” or “issuer” risk and “counterparty” risk. When this happens, “investment” credit risk is normally defined as excluding the holding of deposits, precisely because the issuers of deposits (banks) may also be counterparties to transactions. This can be a natural dividing line if the central bank invests only in very high-quality securities (eg effectively default-free sovereigns), so that there is no real overlap between the categories. As more of the portfolio becomes subject to significant credit risk, though, the rationale for this strict separation loses strength.

Partly reflecting different practices, there is a lack of consensus on the definition of “counterparty” risk. Often, but not always, the term “investment” risk is limited to the holding of primary securities, thereby excluding deposits and derivatives, which are treated as part of counterparty risk. Sometimes, in the context of FX reserve management, central banks define counterparty credit as including settlement exposures, sometimes not.¹

Given these difficulties, the following discussion tries to avoid using specific definitions of “counterparty” and “investment” credit risk. Rather, it breaks down the various types of exposures into their constituent components to see how they are managed.

Measurement of exposures to individual obligors

Individual exposures to obligors can be measured using different metrics. The simplest one is to use just nominal amounts. These are the natural ones for the risk of capital loss in settlements and, possibly, short-term claims that can be redeemed at par, such as deposits. It is also a natural metric if credit risk is modelled in default, rather than mark to market mode, ie if deteriorations in credit quality short of default over the relevant horizon are not taken into account. Another possibility is to use (current) market values. These are the relevant measure when considering the current replacement cost of contracts, as in derivatives transactions or whenever a lag between trade and settlement exists. A third possibility is to also include an estimate of the potential change in the replacement cost arising from changes in market prices (“potential future exposures”).

Whenever different exposures to a given obligor are measured using different metrics, an issue of aggregation arises. One possibility is simply to measure subexposures separately, without the need to synthesise them into a single figure. This is associated with limit systems that either do not cover all exposures in a single limit or where separate sublimits apply (see below). Alternatively, it would be possible to apply weights to the different subcomponents.

Aggregation of exposures across obligors

Aggregation over different obligors can be of at least two types. One is simply to add exposures of the same type or any weighted exposures. Another, more rigorous, one is to take explicitly into account estimates of the correlations of default across obligors. In recent years, advances in risk measurement technology have made this possible, thereby allowing for the calculation of credit risk on a portfolio basis. Even so, the robustness of any such estimates is still subject to doubt. Otherwise, exposures to obligors are simply added, possibly using weights based on instruments and/or measures of credit quality, or by considering subcategories deemed to be comparatively homogeneous.

A complementary approach, which implicitly relies on some form of aggregation, is to carry out stress tests. While possibly varying widely in terms of their specifics, they all have in common the assumption of some “shock” affecting the creditworthiness of parts of the portfolio and/or set of counterparties.

Structure and management of credit limits

The structure of credit limit systems generally follow the risk measurement practices. As regards *individual obligors*, at one end of the spectrum central banks may have a single global system covering all forms of exposure. At the other end, central banks may have separate limits for sets of exposure, possibly excluding some of them altogether. In addition to separate limits for different categories of obligors, some central banks may also have sublimits within an overall total, eg by instrument.² If settlement exposures are not included in limits at all, they may be managed indirectly, such as through limits on the size of the transactions.

The criteria for setting limits for an obligor can vary. Normally, they combine minimum ratings with various types of balance sheet information (eg capitalisation). They may also, for instance, vary across types of counterparties (eg banks and non-banks; OTC versus other) and the maturity of the exposures.

Limits on individual obligors are often complemented with concentration limits at a more aggregate level for the portfolio. These are designed to maintain sufficient diversification. These limits may apply to particular sectors, such as banks, and/or countries.

The enforcement of the limits will depend on the frequency of monitoring and the systems in place. At one end of the spectrum, limit control could be ex ante online; at the other, compliance could be monitored ex post, say on a daily or weekly basis. In the case of ex post monitoring, it is important to take into account intra-period exposures and not only those prevailing at the end of the interval of observation. The ultimate response to the breaches will generally depend on their nature and the context in which they occur.

Other risk mitigation devices

In addition to setting limits on exposures, collateral is one way of reducing credit risk (specifically, the loss-given-default). In FX reserve management, collateral is primarily used for reverse transactions (eg repos, reverse repos, securities lending). One question concerns the criteria for deciding on eligible collateral. Low credit risk and high liquidity are generally important, but there can be different ways of making these concepts operational and different risk tolerance thresholds among central banks. Another question concerns how haircuts are set and managed. While market practice provides an obvious benchmark, central banks may vary in the degree of conservatism.

Integration with domestic monetary policy operations

FX reserve management operations are but one type of operation giving rise to credit risk. Central banks also face such risks in carrying out domestic monetary policy. This raises the question of the degree of integration, or lack thereof, between the two. Because of the different purpose, nature and characteristics of the operations, the integration can be quite limited. But those central banks that take a more balance sheet-wide approach to reserve management and have the necessary systems in place may opt for a more integrated approach, possibly monitoring and enforcing joint limits.

¹ As an example of a possible range of definitions, in the context of estimating exposure at default in the Basel II Framework, the Basel Committee (Annex 4 to the Basel II Framework – June 2006) defines counterparty risk as covering only credit exposures from transactions that give rise to bilateral risk of losses, where the risk of loss arises because of the possibility of default before the final settlement of the transaction cash flows. This would cover the credit risk associated with: (a) derivative exposures; (b) repos/reverse repos/securities lending; (c) the replacement cost risk arising from the lag between trade and settlement dates of “spot” transactions; and (d) the asynchronous nature of some spot transactions (owing to absence of delivery-versus-payment (securities) or payment-versus-payment (FX) mechanisms); (e) the holding of deposits; and (f) the holding of securities/extension of loans. In FX reserve management risk management practices, some may wish to include (d), (e) and/or (f) as part of the definition. ² If settlement exposures are not included in the limits and are not reduced at the source via DVP and PVP or conservative practices in day-to-day timing of disbursements, they may be constrained indirectly through limits on the size of the transactions.