Chapter 1
Why Manage Foreign Exchange Risk?

One of the more puzzling questions facing corporate management is the risk management decision. This is about deciding whether to use a company's resources to actively try to reduce risk, here understood as variability in corporate performance. Our focus in this chapter is therefore to understand when it makes sense for a firm to reduce variability in its performance through managing its FX exposures. Why would a firm be better off managing risk than just accepting its exposures as part of doing business?

To fix our ideas, we can take hedging as a proxy for corporate risk management. Hedging is the use of financial derivatives to alter a firm's exposure to a specific market risk. Many other forms of risk management are possible, like buying insurance, operational risk mitigation, relying on equity rather than debt financing, and keeping a buffer of cash to deal with adverse outcomes. We will in fact emphasize that hedging, while highly flexible, is for the most part only able to address relatively near-term exposures (more on this in Chapter 9). By changing the currency exposure at an operational level (i.e. actively trying to net out cash inflows and cash outflows),

or by borrowing long term in foreign currency, the firm is likely to achieve more durable reductions in FX risk.

Focusing on hedging for now, however, allows us to keep the terminology brief and tap into a rich literature that has developed on why firms should hedge their exposures. Most of the arguments developed in this chapter transfer over to other forms of risk management.

The business case for hedging is far from obvious. In some companies, questions like 'Why should we hedge?' and 'Should we hedge cash flows or earnings?' never seem to get fully settled. The issues are raised, meetings are held, memos are written, and policy documents established. Only for the question to get raised again a couple of years later, prompting new discussions and memos in what seems an endless cycle. A clear and robust answer to these questions continues to elude many companies.

Why is it so difficult to pin down why and when firms should use hedging? We are convinced that part of the difficulty lies in the fact that there is an almost infinite number of performance measures that potentially *could* be hedged. For most of these, a plausible-sounding argument to reduce variability can be developed quite easily, because managers are naturally averse to risk and need little convincing that variability in performance, however measured, is 'bad'. For example, it is not uncommon for business units and investment projects to be evaluated based on their operating margin. For the people involved in these projects, 'protecting' these margins through hedging may seem perfectly logical and desirable. Seeing currency fluctuations eat away at the margins, perhaps even causing them to turn negative, can be a terrifying prospect. The desire to protect margins is behind a large number of hedging proposals that are sent further up the corporate system for management, and sometimes even the board of directors, to decide on. So, what should they say to such proposals? Based on what principles?

In what follows we review the arguments in favour of corporate hedging. The goal is to develop a hedging policy that guides management's decision-making, where a plausible case can be made that the advantages of hedging outweigh the disadvantages.

The alternative to a unified hedging policy, we argue, is an ad-hoc approach at least partly captured by different local agendas, pushing for the hedging of specific business projects, which only serves to make hedging incoherent and unpredictable.

From Individual Risk Management to Corporate Risk Management

In trying to come up with some suggestions for what can constitute a basis for a hedging policy in firms, one might consult the academic thinking on the topic. Academic finance has for several decades been analysing the circumstances in which hedging creates value for a firm.

In the literature, the hedging decision of the firm is sometimes contrasted with that of the individual trader or investor. When we are dealing with an individual it comes down to one thing: risk preferences. Why? Because the situation concerns an individual who makes decisions affecting *his or her own* expected wealth and risk. Therefore, it comes down to what you are comfortable with as an individual. Some people do not like risk and prefer a safer alternative. They are the 'risk-averse' individuals who may prefer to hedge as much as they can. Other people are more tolerant of risk. They may consider it fine to retain certain exposures to risk so long as they also get a reward in the form of a commensurate upside potential. Still others may be positively 'risk-loving'. They are the kind of people willing to take significant risks for the pleasure of the thrill.

The broader point here is that for individuals, the hedging decision follows risk preferences. Once the odds and the range of potential outcomes are known, it all depends on this parameter.

Unfortunately, the model for decision-making based on personal risk preferences does not carry over to firms. This is true despite the fact that businesses are run by people. One might think that risk-averse managers will hedge more and risk-seeking ones less, and that is it. But this does not provide a robust business case for hedging. Managers are in fact hired by the owners of a firm

to run its operations on a daily basis. In more formal language, they are agents hired by a principal. According to corporation law, the managers and directors of the firm are obliged to use their decision-making powers in the 'best interest of the firm'. This pretty much rules out using managers' own feelings and sentiments about risk as a justification for risk management decisions.

Instead we need to think about how hedging can promote the best interest of the firm. Phrased a bit differently, this is the question of how hedging can increase shareholder value. It changes the question from 'How do we feel about this risk?' to 'How can we demonstrate that shareholders are expected to be better off with this hedge than without it?' This is a much trickier question to answer.

The Case for FXRM: Cash Flow Hedging

The basic thing achieved by a hedge, if correctly executed, is to reduce variability in corporate performance. Let us for now take performance to mean cash flow. This is the metric favoured by academics due to its close relation to firm value, which is given by the discounted value of future free cash flow. Reducing variability in cash flow by itself falls short, however, of presenting a convincing business case for hedging. Why? The answer lies in the symmetric impact of hedging on cash flow variability. Say that we enter an FX hedge, which lowers our maximum loss by US$20 mn. The symmetry of a forward contract implies that we also lower our maximum gain by US$20 mn. A firm that protects itself against unfavourable scenarios has to accept that it misses out on the windfall it would have received in more favourable scenarios. The expected value of such a hedge is zero (or even negative, if there are transaction costs). So, while reducing variability might be 'nice' to some of the managers involved, this is not enough to build a case for hedging.

We can conclude, then, that simply achieving a narrower range of potential outcomes for cash flow does not justify corporate risk management. Instead, theorists have sought to build the case by looking for asymmetries in the way corporate performance is

affected by hedging. Hedging can be valuable when there is a negative consequence from cash flow falling below a certain threshold *and* there is no corresponding positive side-effect at high levels of cash flow. Hedging, by reducing the probability of cash flow dropping below the threshold level, reduces the expected cost of these negative consequences. This, in a nutshell, is the idea behind corporate hedging in academic models.

What are the negative consequences that could happen in low-cash-flow scenarios? It is straightforward. Firms need money to meet important cash commitments, like investing in new projects and paying interest and instalments on loans. There are many other commitments to various stakeholders that also require liquidity, like dividends and pension obligations. When a firm does not generate enough cash flow internally, it may not be able to service these cash commitments. Falling below this threshold usually implies some sort of negative consequence. If the firm cannot invest optimally, it will end up being less competitive. If it defaults on its interest payments, it will be sent into bankruptcy. If it sells assets in a fire sale, it may have to accept a price below the assets' fair value. If it cuts back the dividend, shareholders will grumble and reassess the firm's future prospects.

We have now identified the asymmetry we needed. The firm suffers various negative consequences when it fails to reach a certain threshold level of cash flow, but there are no positive consequences, beyond the pure monetary gain, when cash flow is higher than expected. In the good scenarios the surplus cash flows will simply accumulate as more cash.

Someone might reply that it does not matter if cash flow falls below the threshold level because the firm can always borrow to cover up the difference. There might be a credit line it can draw on, or it might have unused borrowing capacity to mitigate the consequences of cash flow shortfalls. Such untapped resources will obviously have to be considered. But all they really do is move the threshold level of performance. With spare debt capacity and cash buffers the firm can tolerate larger cash flow shortfalls. But at some point, the opportunities to get external funding at a reasonable cost

will be used up. Creditors will worry about the risks involved, and may suspect that managers will engage in excessive risk-taking. These things push up the cost of debt, and financially weak firms may find themselves shut out of the credit markets altogether. Equity financing provides no easy escape, as the issuing costs are substantial and investors demand a large risk premium. In these circumstances, the availability of operating cash flow plays a crucial role in terms of being able to execute the business plan and meet various cash commitments.

To summarize, the academic literature recommends looking for different threshold levels of cash flow performance where different negative consequences would be triggered. This is a test to which all hedging proposals should be subjected. The question to be asked is: 'Would implementing this hedge, in a meaningful way, reduce the probability that cash flow ends up below the threshold level given by our cash commitments?' This question raises the bar significantly compared to hedging to protect operating margins, or simply locking in a 'good' FX rate.

The Case for FXRM: Protecting Financial Ratios

In academic finance, cash flow has the highest pedigree because of its close relationship with firm value. However, corporate performance is more multifaceted in real life. Realistically, more things will matter than cash flow. Once we acknowledge this, things get more subtle and subjective, but the arguments are worth taking seriously.

One way to motivate FXRM is to say that a particular financial ratio needs to be stabilized, or 'protected'. Some firms have a stated objective that, for example, the debt-to-equity ratio should not go above 0.7. Since both debt and equity are exposed to FX translation gains and losses, it may decide to manage its composition of assets and liabilities in such a way that the probability of exceeding this target is minimized. This is an argument for FXRM that is not uncommon in practice. Granted, maintaining a set of target ratios

over time gives a sense of direction and an impression that management is on top of things. But what is there to say that defending a particular ratio actually translates into higher shareholder value? Being a corporate objective by itself does not justify using resources to protect it. An objective is a fairly arbitrary thing, and may not always be connected to shareholder value in an obvious way.

Defining threshold levels of performance in terms of financial ratios can be more easily motivated if breaching them is clearly connected to a negative consequence. Again, we are looking for some asymmetry in the outcome distribution. One way this could happen is if the firm has agreed to a covenant as part of its loan agreement. A covenant is a way for a bank to monitor the firm and decrease the risk of default. If breached, the covenant gives the bank increased legal powers to demand specific actions, or even take control of the firm's policies. Most covenants are financial (i.e. they make reference to financial ratios), and they are usually explicit about the threshold level at which the covenant would be considered violated. This allows firms with financial covenants a clear reference point to use in FXRM.

Another case of financial ratios-based risk management concerns the risk of being downgraded by credit rating agencies. Defending a specific rating is high on the agenda in some companies. Managers may consider it likely that the firm's rating will be downgraded if a particular ratio were to be breached. Hedging can in some cases be used to avoid scenarios involving such breaches of key financial ratios. The major credit rating agencies publish a great deal of information about average ratios associated with their different rating levels. Firms can infer from such tables that level of performance at which they would be at risk of losing their desired rating. For example, according to Moody's rating agency, the average ratio of Funds-From-Operations to Debt of an A-rated company is 34.1%.[1] If a firm wishes to defend its A-rating it should probably not go below this level for any extended period of time.

[1] As per the report *Moody's Financial Metrics™ Key Ratios by Rating and Industry for Global Non-Financial Corporates: December 2016* (Moody's, 2017).

Financial ratios therefore provide firms with a potential reference point for understanding what constitutes a critical level of performance. Exchange rates are, as we shall discuss in later chapters, capable of having a substantial impact on financial ratios, not just because of how they affect operating performance but also due to the translation of assets and liabilities from foreign to home currency.

The connection to shareholder value, however, is less obvious for financial ratios than cash flow shortfalls. It is clear that the firm's creditors will applaud any form of risk management that reduces downside risk because it makes debt safer. But maximizing value for debt-holders is not the purpose of corporate management. In fact, due to limited liability, equity-holders enjoy protection from losses beyond the sizeof their initial investment (i.e. they can simply walk away in a bankruptcy). This feature means that shareholders, at least in certain specific circumstances, take a very different view on risk-taking than creditors. They may even have incentives to bet the house by massively increasing risk (referred to in the literature as 'risk-shifting'). So perhaps defending financial ratios at all costs is mostly doing creditors a favour, but shareholders a disservice?

In most cases, however, it would seem that the interests of shareholders and debt-holders coincide with respect to financial ratios-based FXRM. It is usually a bad outcome also for shareholders if a covenant is violated and banks increase their control over operating policies, for example, by making demands that force management to cut back investments or issue new equity. If a rating downgrade occurs, it may reflect badly on management and increase the firm's borrowing costs. Higher borrowing costs in turn translate into lower shareholder value. The kind of situation where the owners would like to gamble and take a huge risk, knowing that they enjoy limited liability, are, though popular in academic models, a rare occurrence in real life. For the most part shareholders want to keep the firm going and enjoy good relations with the firm's creditors.

Based on the above, we conclude that managing the firm's FX exposures so as to lower the probability of breaching important

financial ratio thresholds can many times be justified. It is another matter if management decides to adjust all sorts of policies in order to defend a rating. If the firm purposefully avoids making investments and acquisitions that it thinks are value-creating for the sole purpose of maintaining a certain rating, it could indeed be doing shareholders a disservice. But this is a bigger issue that falls outside FXRM, so we do not pursue it further here.

The Case for FXRM: Managing Net Income Variability

Next, we consider a third dimension of corporate performance, namely earnings-based measures. We may also speak of 'margin-based' measures more generally as firms keep track of their income on several levels: gross profit (revenue less cost of sales), operating profit (gross profit less operating expenses), and various adjusted measures of 'core profits', to give some examples. The firm's ultimate bottom line is of course net income (also called earnings), the measure of profit and loss communicated at great fanfare to the rest of the world every quarter. Net income often attracts massive interest and commentary in the press, and therefore naturally becomes a key concern for management teams.

Does managing variability in net income present a case for FXRM? This is a complicated question that has not been fully resolved. Two camps seem to exist. One argues that yes, protecting against 'bad outcomes' for net income is a valid goal, because it is so important and followed so closely by analysts and investors. If the firm delivers a poor result, or fails to meet expectations, it will immediately get punished and receive a lower valuation. The second camp argues that no, net income is such an arbitrary and accounting-based measure anyway that it does not make sense to make it the centrepiece of a risk management effort. Add to this the fact that management can greatly influence the numbers through the assumptions they are allowed to make in the process of preparing financial statements. Analysts only care about the core

operating profits anyway, so any gains and losses from derivatives are just 'noise' that they adjust away.

We argue that it is possible to find instances supporting both views, but that the bulk of the evidence supports the latter. There is indeed a rock-hard focus on underlying operating performance. Recently, for example, a Scandinavian large-cap company reported a foreign-exchange-induced loss in its financial expenses that was almost four times the size of its operating profit. While this appeared to be a serious situation, a striking thing about it was that it was received by analysts with what seemed like indifference. As they are trained to do in business schools, the analysts move such gains and losses out of their preferred measure of operating performance, considering them temporary effects that are not informative about the firm's future prospects.

Does the argument that the firm gets punished for large unexpected losses, and that FXRM can help avoid this, really have any traction? It is actually fully to be expected that the market rethinks its view about a company when it reports a larger-than-expected loss. In light of the newly published quarterly report, the market will reconsider the firm's outlook because the report contains information the market previously did not have. This is in fact normal procedure and how reasonably efficient markets are supposed to operate.

The key question is *how* the market prices in the new information. Let's say the firm's income before financial items is US$200 mn lower than anticipated but there is also a gain on an FX derivative of US$100 mn offsetting the lower operating performance. Will the market price the firm according to a multiple of its operating income or its net income? The latter ratio is the famous price-to-earnings multiple, probably the most widely used multiple out there. If market participants price the firm as a multiple of net income, the derivative gain will get capitalized into the price by the same multiple as operating income. In that case hedging will indeed stabilize the value of the firm. But it is not reasonable that the market prices the firm's earnings in this way. Capitalizing the derivative gain with the same multiple implies that it is expected to generate

the same gain in all quarters going forward, which is clearly not realistic. The effect of a derivative is limited to its expiry date, and the present value of *future* hedging transactions is zero. Instead, it is much more likely that analysts see this as a temporary effect that does not affect the firm's future prospects. That is, it prices the firm according to its view of underlying operating performance.

Proponents of earnings-based FXRM to protect and stabilize net income need to rely on subtler arguments. Perhaps investors just do not 'like' net income volatility on a more general level? If so, they will pay a premium on stocks whose earnings are more stable, even though the expected mean is the same. The claim that there is a risk premium on net income variability is precarious in light of existing theory about how assets are priced. But many managers seem to believe it nonetheless, as evidenced by the extent to which they engage in 'earnings management', that is to say making convenient assumptions regarding the value of its assets and liabilities to smooth net income. A legendary case in point is General Electric, who under Jack Welsh's leadership presented growing net income from continuing operations in 100 consecutive quarters.[2] Throughout this period, core earnings (i.e. net income adjusted for managerial discretion) were lower and much bumpier. According to many academics in finance, earnings management is based on 'market myths' (i.e. a false and simplistic belief about how the market prices the firm's stock).

The case for earnings-based FXRM thus rests on shaky academic grounds. But it is hard to entirely discount the possibility that in some situations it makes sense to reduce variability in net income. The crucial issue as far as FXRM is concerned is whether different external stakeholders make decisions that affect the firm based on its net income performance, and whether those decisions can be swayed in a favourable direction by FXRM.

As before, we need to look for negative and asymmetric consequences of net income volatility to defend net-income-based risk management. Perhaps the firm is under a lot of pressure to

[2] See Roger Lowenstein's book *Origins of the Crash: The Great Bubble and its Undoing* (Penguin Books, 2004) for a full account.

improve its performance and analysts are losing their faith in the management team's ability to run the business. Perhaps also its financiers are paying a lot of attention to how its net income develops, and that their willingness to provide additional funding depends on net income not falling below some reference level. This could be a very unwelcome time for the firm to experience FX-related losses. FXRM could then be valuable to the extent that it reduces the risk that net income falls below the perceived threshold value where the investors would conclude that the firm is a hopeless case and withdraw their financing. While this setting is not entirely inconceivable, it does depend on several fairly strong assumptions, like investors following simplistic decision rules with regard to net income performance. The burden of proof in cases like these should lie on the individual(s) who argue in favour of hedging on these grounds.

Key Chapter Takeaways

- Risk management (i.e. reducing variability in corporate performance) does not automatically translate into higher shareholder value.
- Managers' dislike of risk and performance variability does not constitute a sufficient business case for FXRM.
- Academic models of value-enhancing risk management are built on the idea that there is an asymmetry: poor performance triggers a negative and costly consequence, but good performance only leads to more cash accumulating.
- Negative consequences of cash flow variability include not being able to invest optimally or maintain dividends; a failure to meet interest payments; or having to make asset fire sales.
- Managing the risk that cash flow falls short of important cash commitments presents the strongest argument in favour of FXRM.

- ➢ FXRM can also be justified on the grounds that it makes breaching key performance thresholds for financial ratios less likely, for example, related to loan covenants or credit ratings.
- ➢ Stabilizing earnings, or some other profit margin, tends to be the weakest argument for FXRM and must be critically assessed before being accepted.

Further Reading

Froot, K. A., D. S. Scharfstein, and J. C. Stein. 1994. A framework for risk management. *Harvard Business Review* **72**, 91–102.

Smith, C. W. and R. M. Stulz. 1985. The determinants of firms' hedging policies. *Journal of Financial and Quantitative Analysis* **20**, 391–405.