use, such as sight, sound, touch and smell. So as we look at Vanbrugh's architecture or listen to Telemann's music we are continually but unconsciously taking apart its constituent elements and predicting how it will behave next. In fact, without such a process the world around us would simply generate too much information for us to deal with. In the last chapter we saw how the notions of attention and foreground and background help us with this problem. However we also rely upon structured and ordered rule-based events in order to survive this massive bombardment of information. To take it to the extreme, a totally random sequence of sounds does not seem to be music at all. Even the musical scale with its restricted range of notes organized in strict ratio of pitch is a huge structuring of the otherwise infinite range of possibilities. Whilst modern composers may experiment with atonal music and the removal of timbre, it is highly unlikely ever to become popular.

Order, pattern and redundancy

Armed with this view of human perception as active and predictive, we can now see that as an experience unfolds we can examine the extent to which we have prior knowledge about it based on our ability to predict it. Let us consider a very simple model in order to explore the way this works. We are watching a coin being tossed. We know from our previous experience of this kind of event that on about half the occasions this will come down heads and the other half will be tails. Both outcomes are equally likely, and no other result is possible. (Strictly speaking the coin could land on its edge and miraculously stay there, but this is so unlikely that we can discount it!) So we have some advanced knowledge about the outcome of this event, and if we guess we are likely to be correct about half the time. The actual taking place of the event therefore does not remove all the uncertainty it would if we had no prior knowledge. Similarly, if we are watching a dice being rolled and trying to guess the outcome, we also have some prior knowledge, but rather less. Here only about one in six of our guesses is likely to be correct. The event itself then removes more uncertainty than does the tossing of a coin.

Let us make the scenario a little more realistic and complex. We are waiting at a bus stop used by four bus routes. Again we might guess as to which route the next bus will be on, and we would be right on average one-quarter of the time. Perhaps, though, buses on these routes are not all equally frequent – the first route might have buses running every five minutes, with the second route running buses every fifteen minutes, and the other two running every half-an-hour. If we had no knowledge of this pattern of frequency we could not use it to help us guess. However, regular users of this bus stop would be aware of the timetable and could use this to raise their rate of correct guessing to