prescriptive rules based on theories and fantasies far removed from the serious business of predictive underwriting.

The road to automation

But back to the year 2000. I had to convince the board, and then the rest of the business, that full automation was the way forward. There were many discussions and even more sceptical colleagues.

It was more straightforward to convince the credit team of the benefits of automation. They could immediately see the advantages in terms of increased quality. The main issues arose with operations and sales. There was a subconscious desire to hold onto the traditional way of doing things.

Operations staff in particular sometimes wanted the new system to do old jobs. The sales team was worried about how intermediaries would react to such a "black and white" underwriting approach, without influence or leverage.

The board was persuaded, and it empowered our credit team to re-engineer our underwriting process. So the team set about specifying an automated process, and looking for a supplier. Team members considered a number of routes forward. Building our own bespoke system from scratch, using different suppliers for various different components, was a strong contender.

However, GMAC-RFC had already introduced in the US a decision-making engine called Assetwise, for which our colleagues were seeking a European showcase. The system worked differently in the US in that applicants' details were keyed in and the computer found the best terms for that particular applicant amongst the GMAC-RFC range. As it is such a commoditised mortgage market in the US, intermediaries and their clients were interested in understanding the rate they would pay at a particular moment in time.

In the UK, we felt that demand would be different. Intermediaries largely know, and have sold to the customer, the product they will be offering before applicants' credit details are entered. So we needed to convert the system to UK needs and re-write to our specification, which turned out to be a longer lead-in time than we had initially planned. The system we ended up with, however, was certainly something special.

By the end of the year we had not only completed our outline systems specification, but had chosen to go with Assetwise. We started work on the system in early 2001 and successfully delivered automated decision-making at Bracknell fairly quickly - the first step to our revolutionary Point of Sale Decisioning system. This enabled our internal underwriting to be slicker and quicker with reduced references. But we then hit a roadblock with our plans to extend the system to the point of sale.

When we had purchased the business on which our UK launch was based, legacy computer systems had been involved. It turned out that these were too old to accept either the new technology we were proposing to bolt on in terms of the Assetwise decision engine, or the volumes we were proposing to pass through them. To simplify a complex situation, we had to re-write the front end of the system before we could make real progress towards POSD. From one perspective, having waited so many years to have the opportunity to fulfil this part of my vision, it was relatively easy to be patient and wait another year or so. Already the automation was delivering interim wins in terms of making the burden of processing lighter.

Our volume history does not show that this delay was detrimental to our growth ambitions. On the other hand, it did allow other lenders to catch up and, eventually, overtake us in terms of launching online decision making , although most of the lenders in question eventually launched what I would describe as Decisions in Principle.

While we undertook this systems re-write we took the opportunity internally to make decisions on all of our products through the new engine, starting with mainstream and then adding self certification and, finally, non-conforming. Many felt, and some still feel, that automated decision making is best placed to underwrite prime or niche business only. But we have found that it is equally predictive for non-conforming lending where we know the applicant's credit history is not perfect and can use this to accurately risk-price and determine future performance.

By the autumn of 2002 we were back on track following the systems re-write, and in early 2003 able at last to introduce a select number of intermediaries to the pilot testing phase of our new point of sale decisioning engine. We tested the impact on intermediaries who were not regular supporters, and intermediaries who were.

The most interesting data came from the former category, who started to use the system for much more of their business, such was its ease and simplicity. We had reduced to an average of just 14 questions the information that needed to be keyed in order to get an immediate online decision, and we were maintaining response times around about the 30-40 seconds mark.

We knew it would be important to keep the questions to a minimum, and to make them dynamic. Some lenders make intermediaries go from page to page on the screen, often collecting data that is irrelevant to that particular case, and leaving the intermediaries to wonder when they are going to come to the end of the process.

When an accept decision was delivered we allowed the intermediary to stay online and book funds as well as instruct the valuation on payment by credit card. Often this meant that the valuation report came in before the application form.

We ensured that training was proactive and clear, meaning that we quickly got to 95% usage without the need for the financial incentives that other lenders had introduced when launching their online systems. Our focus from day one on intermediary ease of use had really paid off.

Some performance statistics

The next big stage of the testing process was to eliminate referrals. Again, this was a piece of contra-thinking. We were not going down the same path as other lenders and maintaining the manual underwriting process alongside an automated one. We were shooting for the bigger prize of no doubts, no manual intervention and total certainty. We called it “definitely, not maybe”. It meant therefore that we had to eliminate referrals, where the system does not reach a decision and, instead, refers the application for manual underwriting. We had been working on this in parallel to the systems re-write. Figure 4.1 shows our progress in the area of mainstream loans in getting referrals down from an initial 46% to nil by the time we launched to an expectant market in June 2004.

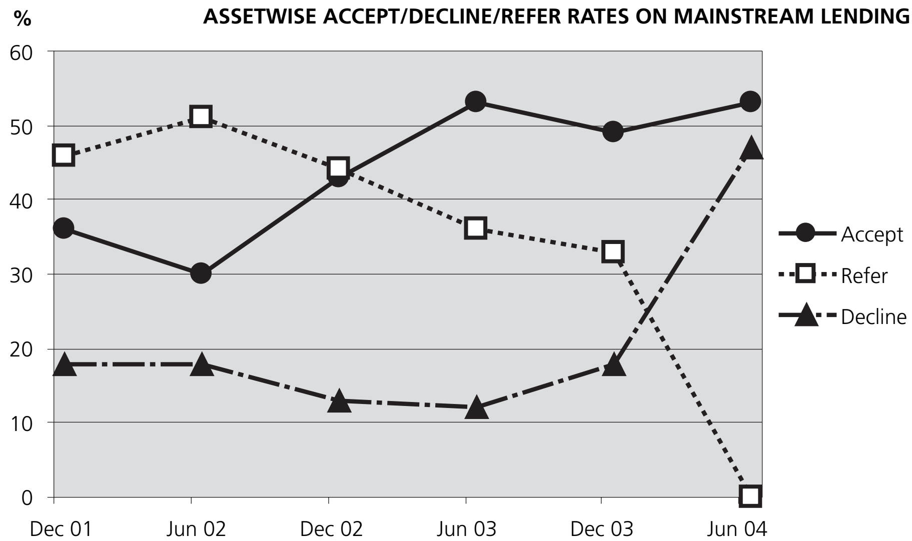

Figure 4.1

The vast majority of "referrals" were longwinded, complicated, frustrating "NOs". A minority were longwinded, complicated, frustrating "YESs". Our Unique Sales Proposition was that we would be the only lender offering binding credit decisions at point of sale without any referrals. Decisions that intermediaries could rely on and use for planning. We therefore had no choice but to eliminate the referrals, and this is what we did. (We do allow a limited appeals process for those few "no" decisions that prove to be the exception to the rule).

There are, of course, those who take a contrary view to ours. Some lenders feel it is a service to intermediaries to deal manually with all referrals. My personal experience in running Private Label and GMAC-RFC, is that referrals rarely do anybody a service. The majority still turn out to be declines after a lot of paper is gathered and emotional energy expended - time that could be spent more efficiently processing better quality applications.

By offering a fast, automated, easy decision we were allowing intermediaries to know exactly where they stood. By not having manual referrals, we would be able to deliver cost savings that could be reflected in the price at which we could make mortgages available. We would also, I was sure, increase the quality of our lending over time through the elimination of manual intervention.

We launched POSD in June 2004, a couple of years later than we had intended, but well worth waiting for. Prior to the introduction of automated decisioning, we prided ourselves on the fact that we were able to give intermediaries a decision in principle on a complete application within 25 minutes. After the introduction of our POSD system, this timeframe reduced to 30 seconds on average.

Moreover, the decisions we were able to give were binding, subject only to valuation and the non-discovery of fraud, as opposed to the DIPs available under manual underwriting. Figure 4.2 shows that dramatic fall in more detail:

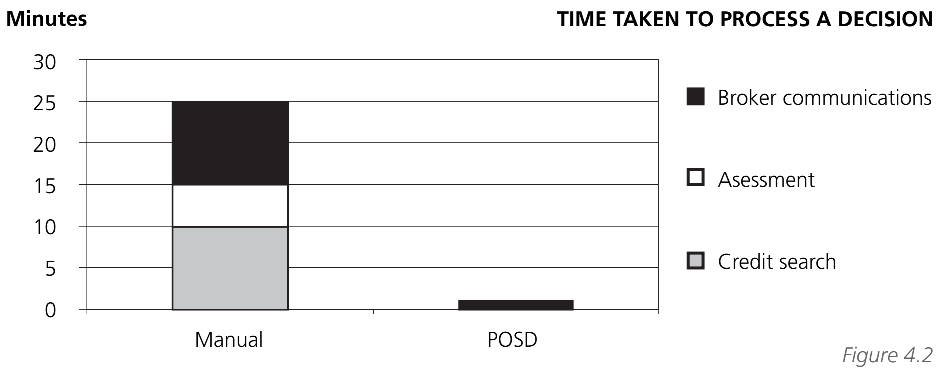

Figure 4.2

We saved about 50,000 hours of work in the first year of POSD related solely to automated vs manual decisioning. No lender, of course, receives a formal application for every decision made so this is reflected in the calculation.

But the benefits of POSD did not stop there. After a DIP had been given in the pre-automated manual world, the resultant application had to be processed. We thought we were being slick in getting an application on average from receipt to completion in two and a half hours of exclusive work on that case, according to our capacity planning model at the time. But in a POSD world the average processing time per application reduced to just 40 minutes, as shown in figure 4.3:

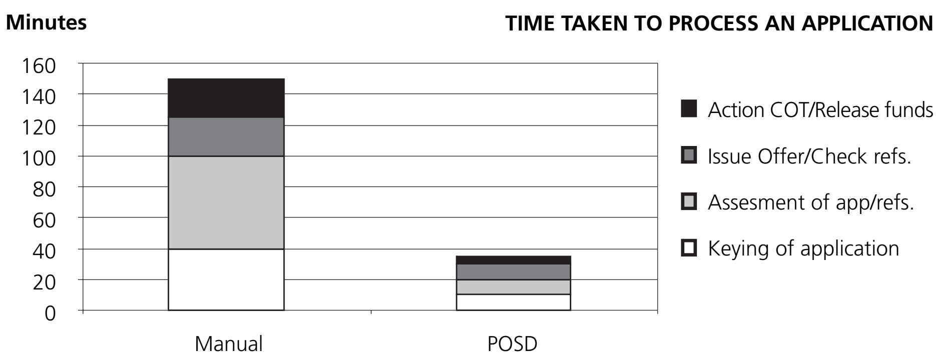

Figure 4.3