National Audit Office

Financial Ombudsman Service

Report on the model of the case handling system

Technical Appendix

MARCH 2012

Contents

Part One

Introduction

1.1 The Financial Ombudsman Service (the Service) was created by the Financial Services and Markets Act 2000 as a scheme under which certain disputes between consumers and financial services institutions may be resolved quickly and with minimum formality by an independent person. Its jurisdiction was later extended by section 59 of the Consumer Credit Act 2000.

1.2 The Board of the Service have commissioned the National Audit Office to perform a review of the organisation, with a particular focus on efficiency. One of the areas under review was a new tool the Service developed to plan resource capacity based on forecast demand. The tool's development was part of the Service's Demand, Capacity and Operational Planning project (DC&OP tool). We needed assurance on how the tool has been used to make a planning decision, and test its accuracy.

1.3 To do this we developed a simulation model to reproduce the end-to-end case handling process, to draw conclusion on the sensitivity of the process to volatile demand.

The Service pathway

1.4 We developed the case handling pathway model by reviewing the Service process maps, carrying out a process walkthrough and interviewing process owners. The logic flow map is shown in Figure 1.

1.5 The case handling process has three main stages: customer contact division stage; adjudicator stage; and decision stage. For the purpose of our analysis, we divided the adjudicator stage into preparation and adjudication sub-stages, and the decision stage into preparation and final decision sub-stages.

Customer Contact Division Stage (CCD)

1.6 CCD is the first point of contact between the customers and the Service. CCD staff provide explanation of the process to consumers and alert businesses that they have been contacted. When enquiries are not in scope, consumers are advised where to take the case. If the consumer is dissatisfied with the business's response to the complaint, or if the business has not responded to the complaint within eight weeks of first receiving it, CCD will convert it into a case.

Preparation for Adjudication

1.7 Once a case has been converted, the consumer and the business are required to submit the supporting evidence and documents. Staff within the operations support pool are responsible for processing the incoming documents and updating the case files.

Adjudication

1.8 Cases are then allocated to adjudicators who check jurisdiction, gather evidence from the business and consumer and recommend a resolution. If both parties agree, the case is resolved. The majority of cases are resolved at this stage. Alternatively, either party can ask for the case to be reviewed by an ombudsman. If so, the parties are invited to submit any further evidence they wish to be taken into account.

Preparation for Decision

1.9 The ombudsman support team check that the case file has all the information that will be required by the ombudsman. A small proportion of cases are resolved at this stage. The support team will prepare a possible draft decision for the ombudsman to consider, if the adjudicator has not already done so in a process referred to as self-prepare.

Final Decision

1.10 The Ombudsman reviews the case file and writes an independent decision on the outcome. If the Ombudsman's view is significantly different from the adjudicator's recommendation, the ombudsman may issue a provisional decision first, to give the parties another opportunity to comment. The Ombudsman's final decision, if accepted by the consumer, is legally binding on both parties.

Re-work and re-open cases

1.11 If more evidence is needed at the draft and final decision stages, the case files are sent back to the adjudicator to contact the relevant parties and make the requests.

1.12 Depending on the circumstances, and unless a final decision has been issued, it may be possible to reopen a case if new evidence has been found.

Figure 1

Logic map of the NAO's model of the case-handling process

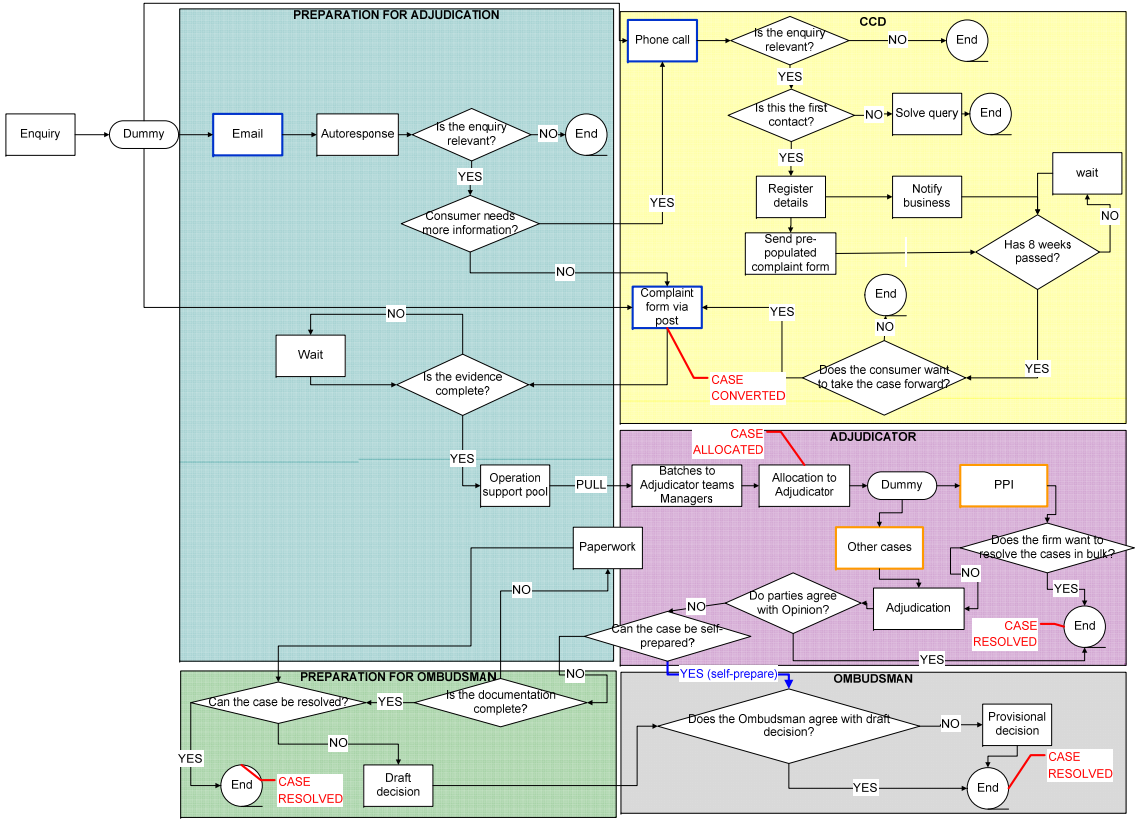

Source: National Audit Office

This logic map was developed by the National Audit Office study team following observations of the process and interviews. The logic map was developed in order to inform the development of the model.

The Service case mix

1.13 In 2010-11 the Service received just over 1 million initial contacts, of which 206,121 were converted into cases.

1.14 The Service handles complaints for a varied mix of cases. Cases are analysed according to three main product families (banking, insurance and investment) which are subdivided into 18 executive product groups. Because of the surge in complaints about payment protection insurance (PPI), that is currently analysed separately from other insurance products, and analysed into six subgroups based on the size of firm and whether the complaint came through a claims management company (CMC). See Figure 2.

1.15 Cases across the different groups follow the same process within the Service, and are allocated to staff depending on their area of expertise. The time necessary to resolve cases varies accordingly to the type of case and the case's complexity.

Figure 2

Product families

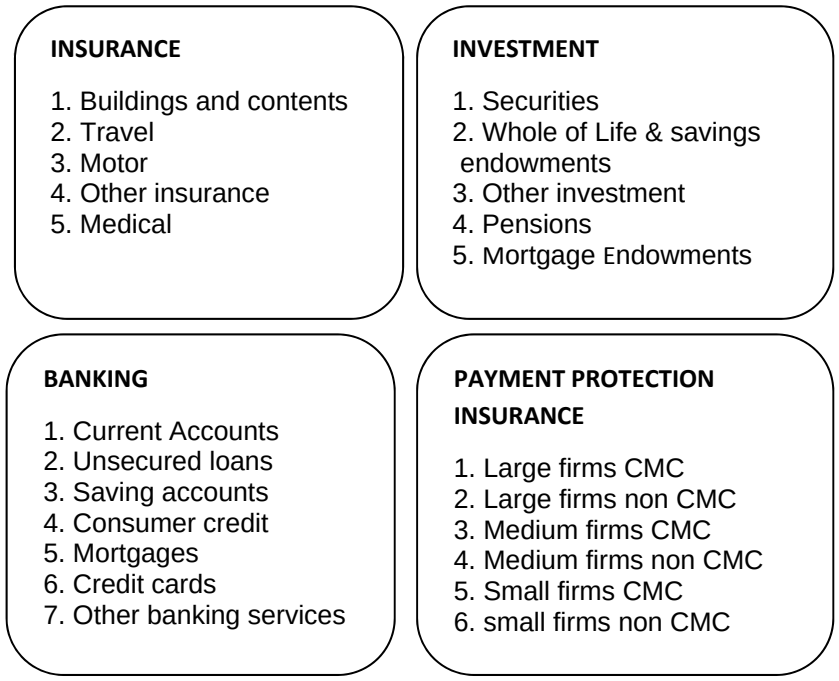

NOTES

Source: National Audit Office

Part Two

Modelling method

2.1 We used discrete event simulation to develop a visual interactive model of the Service's operations. The model simulates how cases flow through the end-to-end process over time, from case set up to final resolution.

2.2 The inputs for the model were based on the cases processed within a one year period and comprised the elapsed time for processing individual cases and the number of resources used. The elapsed time for processing individual cases was divided into the time spent completing each of the stages in the process (cycle times). The model enabled us to:

- better understand the behaviour of the system on a two years horizon;
- monitor the time spent in the system by single cases;
- gain insight on where problems might exist; and
- explore modifications to the system.

2.3 By doing this we were able to predict the effect of a change in demand for any one type of case (with the assumption that the complexity remains constant) on the time it takes to progress, the length of the queues and what the effect of changes in resourcing will be.

2.4 We developed the model using the simulation software Simul8, and coding the syntax in Visual Logic, when needed.

Model description

2.5 Given that the scope of the model is assessing the impact of volatile demand rather than identifying solutions to streamline processes, we decided to model the main stages of the process instead of single tasks carried out by the staff.

2.6 While the model structure can be applied to the whole case mix, we decided to focus the analysis on a single product type: mortgage cases.

2.7 To reduce the complexity of the system and account for a lack of data, further assumptions and simplifications were made:

- a year has 260 working days;
- cases can be reopened even if closed at adjudication stage (see paragraph 1.12);

- these re-opened cases go through the process again from the start (in reality they would go back to the adjudicator); and
- cases are not prioritised (in reality there is a prioritisation system, however there was no data with which to model this).

2.8 In particular, we decided to exclude from the model:

- Customer Contact Division. We modeled exclusively the cases that are set up, and kept out of scope all the enquiries received that do not end up as cases.
- Interactions and communications with consumers and businesses. Whilst this is not explicitly modeled, the time spent liaising with parties is taken into account within the cycle time within each stage.
- The complaints submitted by consumers regarding the way in which the Service has handled their cases.

2.9 The assumptions were discussed with both the Service and a modelling expert. From this we developed the structure of the model (Figure 3).

Figure 3

The model structure

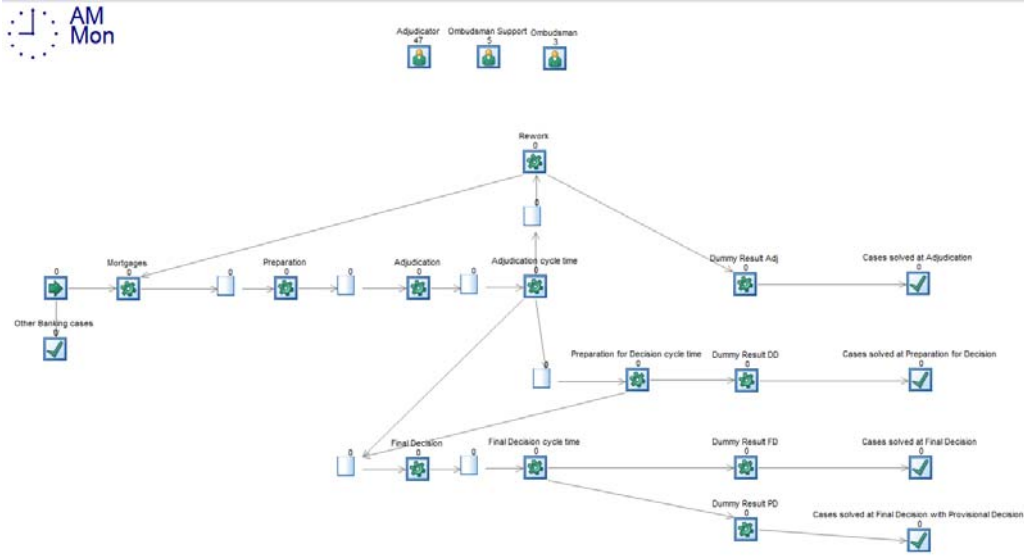

Point at which cases enter the system

Queue - cases wait here until a work centre has the capacity to pick it up

Point at which cases enter the system

Adjudication/Preparation for Decision/Final Decision - work centres which represent the net working time at each stage (see paragraphs 2.23-2.26 for a discussion of this)

Adjudication cycle time/Preparation for decision cycle time/Final Decision cycle time - work centres which represent the elapsed times of each stage (see paragraphs 2.28-2.29 for a discussion of this)

Resources available for net- working time work centres to use (see paragraph 2.27 for a discussion of this)

Points at which cases are resolved.

Source: National Audit Office.

Data analysis

2.10 The model was developed using data on mortgage cases, as received by the Service. When data for the selected product type were not available, we used data of the Banking planning family instead.

2.11 Data was obtained from the following sources:

- To model the probability distributions of cycle times at each stage in the process we used an extract from the Service's case-handling software, Clipper, of all closed cases between August 2010 and August 2011.
- To model the number of resources available we used data from the Service's resource planning tool[1] as at August 2011 (data on the number of resources used by product type were not recorded on a monthly basis before the introduction of the resource planning tool).
- To model the daily volumes of cases we used an extract from the software used by CCD of cases set up between Dec 2010 and August 2011.

2.12 Where data were not readily available, we estimated the parameters according to the expert judgment of the Operations Planning Manager and the Transformation Programme Manager.

[1] This tool was developed as part of the Service's DC&OP project.

Data analysis for label values

2.13 Once cases have arrived in the system, the model assigns certain characteristics to them which then determine their pathway and time spent in the system. These characteristics include:

- the type of case;
- who sent the complaint (i.e. consumer or claims management companies); and
- where the case is to be resolved (i.e. adjudicator, preparation for decision or final decision stage).

2.14 The labels were attributed according to the historical data as extracted from the Service database. Of the banking cases arrived within one year, 10.1 per cent are mortgage cases. Of these, 5 per cent are sent by claims management companies and 95 per cent by consumers.

2.15 We used the labels to route the cases through the system, and assign different time distributions (see paragraph 2.28) according to the stage at which the case is resolved, and the type of case it is (Figure 4).

Figure 4

Label for when a case is to be resolved within the model

Stage	Resolve	Progress	Self-prepare
Adjudicator (%)	81	14	5
Preparation for Decision (%)	23	77	
Final Decision (%)	85	15[2]	

NOTES

The percentage of cases that progress at final decision stage represents the cases for which is issued a provisional decision.

Source: National Audit Office

[2] The model assigns a label to 15% of cases that go on to Final Decision to 'progress'. These cases represent those in which the ombudsman's view has differed significantly from that of the adjudicator's recommendation, and the Ombudsman has issued a provisional decision, to the give the parties another opportunity to comment.

Data analysis for number of cases arrived

2.16 To establish the number of mortgage cases arrived in a day, we used the data collected by the Customer Contact Division centre - the number of banking cases set up in a given week.

2.17 We assumed that the cases were homogenously distributed within the working week, and no cases arrived during the weekends. Hence, we divided the weekly number of cases set up by five to obtain the daily number of cases.

2.18 Using the software StatFit, an add-in for Simul8 to find the best-fit probability distribution on input data, we fitted the number of banking cases arrived in a week to a set of statistical distributions. Due to the nature of the process (a number of events occurring in a fixed interval of time), we expected the distribution to be discrete, and approximate to a Poisson distribution. In fitting the data, we found the best fit to be a Normal distribution with a mean of 1278 cases per week and a standard deviation of 1280, as shown in Figure 5.

Figure 5

Normal distribution fitted for Banking case arrival

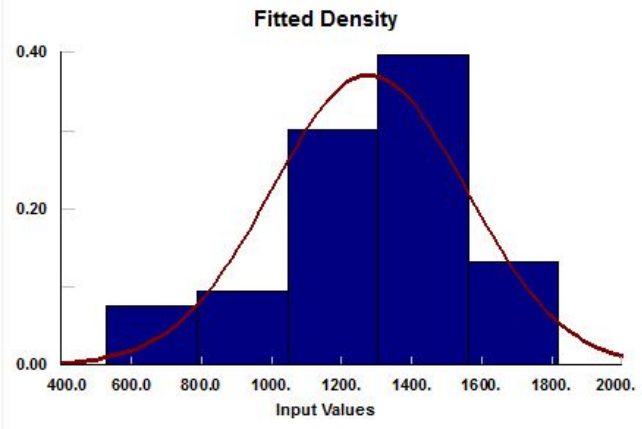

Source: National Audit Office

2.19 Since the mean is a sufficiently large value (>1000), the normal distribution is a good approximation to the Poisson distribution. We hence selected to input in the model the inverse of the Poisson distribution: the Exponential distribution with a mean inter-arrival rate of 0.004[3] days. As mortgage cases represent 10% of banking cases that arrive, the model progressed 10% of banking cases that arrived.

Data analysis for cycle times

2.20 The Service uses Clipper to record a time-stamp at the point in which a case enters and exits a stage in the process, as shown in Figure 6. We used these time stamps to divide the elapsed time into the stages of the process (cycle times). To obtain the cycle time a case spends at each stage we calculated the net working days between time stamps. For example, the cycle time for the Adjudicator stage is given by the difference between the date at which the case is first allocated and, depending on when the case is resolved, the:

- closure date, if the case is resolved at this stage;
- date the case first entered the decision queue for preparation for decision, if the case progresses; or
- date first entered the decision queue for ombudsman, if the case is self-prepared.

2.21 We carried out quality checks on the data, and removed the following:

- cases with negative cycle times, which we assumed were caused by errors in data input;
- cases re-opened or sent back to the Customer Contact Division; and
- cases with cycle times greater than 1000 working days, considered as outliers.

2.22 We also excluded from the analysis of cycle times for the preparation for adjudication stage those cases that took part in a pilot of a triage system which ran for a short period in 2011, in which certain cases bypassed the preparation for adjudication stage and were allocated directly to teams depending on the complexity and expertise of staff.

2.23 Overall, we removed 661 cases from the initial 6217 in the dataset.

[3] 0.004 is obtained as 1/256, where 256 is the mean number of cases received in a given day (1278 times 4.33 weeks in a month, divided by 21.66 working days).

Figure 6

Time stamps

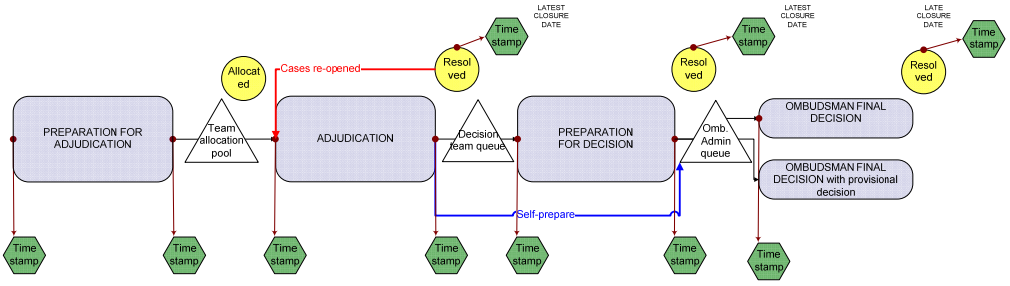

Source: Financial Ombudsman Service

2.24 The calculated cycle times represent the elapsed time of a case at each stage in the process. This includes the time spent by staff actively working on cases, the time spent liaising with parties and the waiting time for additional information. As we did not have detailed information on the break down and durations of single tasks that happen within the elapsed times, we decided to:

- model the preparation for adjudication stage as a 'black box' of infinite capacity. This represents the delay between when the case is initially set-up and when it is ready to be allocated to the adjudicators. It includes both the net working time of staff from operation support, and waiting times; and
- divide each of the remaining stages into two components, as shown in Figure 7.

Net working time

2.25 The first component of these stages represents the net working time spent by a member of staff to process a single mortgage case. The Service provided the monthly net average time spent working (in days) on a single banking case by the different resources. These were 1.66 days for adjudicators, 0.9 days for the ombudsman support team and 0.5 days for ombudsmen.

2.26 We created Erlang distributions with the given mean to generate some randomness in the model, as shown in Figure 8. The Erlang distribution is similar to a Normal distribution, but is slightly skewed to allow for the fact that expected net working time is not symetrically distributed (it cannot be less than zero). We determined the parameter *k* by manipulating the shape of the distribution to obtain some spread with non-null values.

Figure 7

Stage break-down

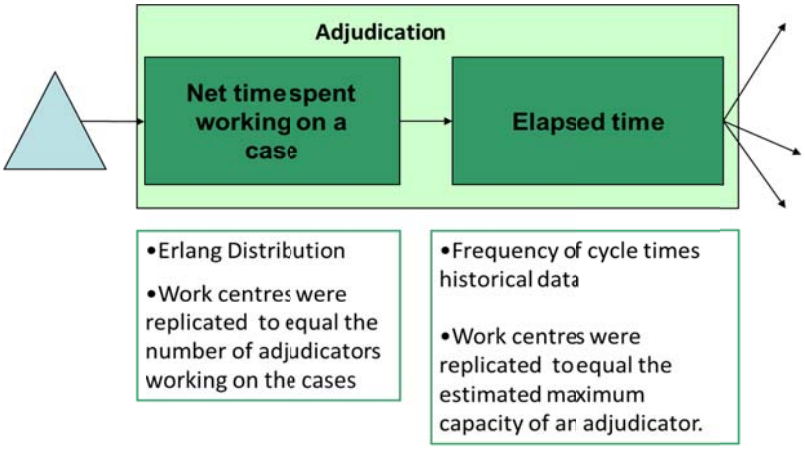

Source: National Audit Office

Figure 8

Erlang distribution for Ombudsman net cycle time

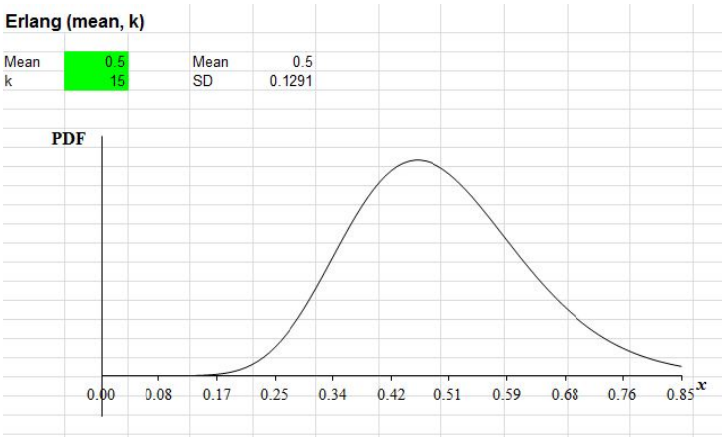

NOTES

1. *K* is a non-negative integer that defines the shape of the distribution.

Source: National Audit Office

2.27 For each stage, we replicated the number of net-working time work centres to equal the number of full time equivalent staff working on mortgage cases.

2.28 Using data from the Service's resource planning tool, we took the total number of days spent working on mortgage cases, and divided by the number of working days in a year to obtain the number of full-time equivalent resources (Figure 9).

Figure 9

Number of full-time equivalent resources

	Adjudicators	Ombudsman Support Team	Ombudsman
Days in the last 12 months	11733	1179	728
Number of equivalent resources	47	5	3

NOTES

1. 22 working days per month.

Source: Resource Planning Tool, August 2011

Elapsed cycle time

2.29 The cycle times calculated from the historical data did not fit any standard statistical distribution. As we did not have additional information to further segment and clean the data, we calculated frequency tables to input the historical distributions of elapsed times directly into the model. We segmented the distributions:

- by type of case, dividing CMC and non CMC cases at the preparation and adjudication stage. This analysis was not possible at the subsequent stages of the process in the model due to the fact that the majority of cases have been resolved, and therefore dividing the cases that do progress produces too few cases to analyse separately; and
- by resolution stage, separating the cases that progress and resolve at each stage.

2.30 We segmented the data in this way since we found that there were statistically significant differences between mean cycle times for CMC and non-CMC cases. In total, we input eight different probability distributions of elapsed times.

2.31 The number of work centres at each stage in the model represents the maximum number of cases that can be worked simultaneously, given a certain level of resources at that stage (see Figure 9). Since data on the system's maximum capacity are not recorded, we estimated the following: 1692 cases at adjudication, 140 at draft decision and 81 at final decision. We obtained the parameters by performing multiple runs and calibrating the model against the historical data.

Model verification

Simulation clock

2.32 The time unit selected for the model is a day. The simulation time is set to five days per week (the Service is open from Monday to Friday), each day starting at 9:00 am and finishing at 5:00 pm (8 hour working day)[4]. Since the average time spent in the system of cases resolved at the final opinion stage is 434 working days (Table 1), we collected the results for two working years.

2.33 We set a 'warm-up' period of 4 years, to allow the system to reach a steady state and operate like the real system with finite resource capacities. Warming up the model enabled us to initialise the queues and work centres that in real life are not empty at the beginning of the year. We established the warm-up period by running the model several times and observing after what period of time the results become stable.

2.34 We used the Simul8 trial calculator to establish the number of times we needed to run the model. The calculator was set up to ensure that the model would be run enough times so that the 95 per cent confidence intervals of the results obtained would be within 5 per cent of the estimate of the mean.

2.35 Including the warm-up and result collection period, the model was run for six years (1560 working days) and eleven replications were made for each scenario to allow for the variation in results.

Conceptual Model Verification

2.36 In the early stages of developing the conceptual model we discussed it with the Service and took into account its comments.

2.37 We undertook a verification process throughout the model-building period. All the label values that are attributed to cases by the model were evaluated to ensure that the model obeyed the coding of variables. As the model progressed, we stopped the simulation and checked that the different cases followed the correct path.

2.38 We tested the logic of the model and the coding with a discrete event simulation expert. We observed how accurately the model outputs compared to the historical data, and observed whether the model correctly reacted to changes in the inputs.

[4] These were the timings used for the model. The actual hours that the Service is open for contacts is from 8am to 6pm Monday to Friday and from 9am to 1pm on Saturdays.

Part Three

Results

3.1 The current scenario results provide an insight into current performance. For the experimental scenarios, we simulated what would happen to the time to resolve cases and to queuing times, when a change in demand occurs, and how the Service could react. When testing alternative inputs (resources and demand), we used the same random numbers for all replications in order to enable comparisons.

Current scenario

3.2 Table 1 shows the results of the model, given the current level of resources and number of cases arrived. The results aligned with those reported by the Service, for the year August 2010 to August 2011. Figure 10 and Figure 11 show how the historical data fall within the interval of the values obtained running the model eleven times. These values were set as baselines of comparisons for the experimental scenarios.

Table 1

Results: current scenario

Resolution stage	Number of cases resolved	Average time spent in system (days)	% cases solved within 18 months	Average queuing time (days)
Adjudication	9673	139	99%	12
Preparation for decision	334	432	41%	160
Final decision	1472	381	54%	15
Final decision with provisional decision	258	434	40%	15

Source: National Audit Office simulation modelling

Figure 10

Comparison between the historical data and the model results for number of cases resolved

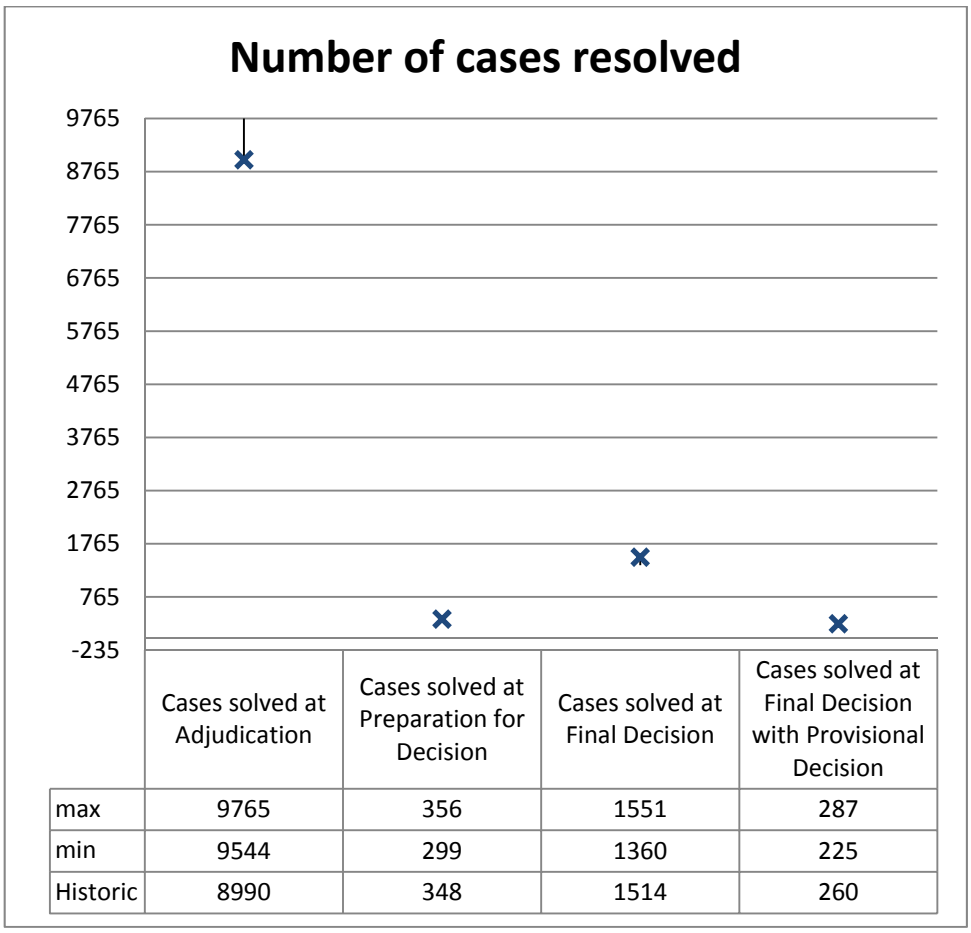

	Cases solved at Adjudication	Cases solved at Preparation for Decision	Cases solved at Final Decision	Cases solved at Final Decision with Provisional Decision
max	9765	356	1551	287
min	9544	299	1360	225
Historic	8990	348	1514	260

Source: National Audit Office

Figure 11

Comparison between the historical data and the model results for average time spent in the system

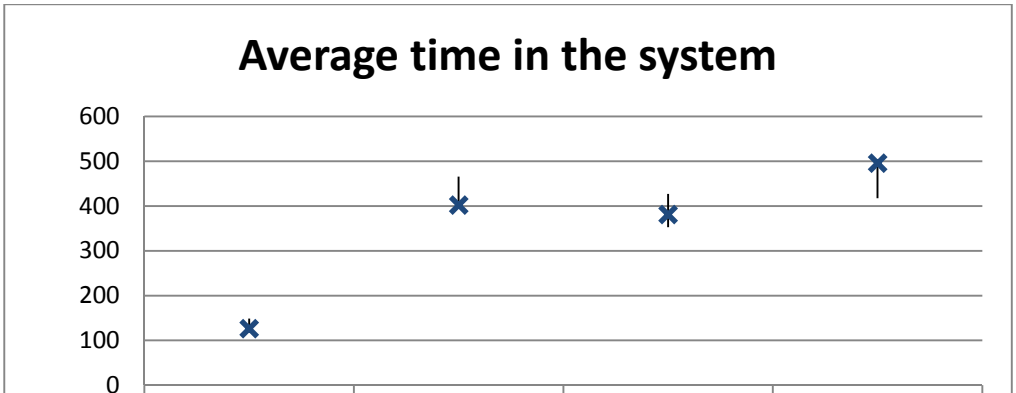

	Cases solved at Adjudication	Cases solved at Preparation for Decision	Cases solved at Final Decision	Cases solved at Final Decision with Provisional Decision
max	149	466	427	466
min	135	395	353	418
Historic	126	402	381	496

Source: National Audit Office

Experimental Scenarios

3.3 In October 2011 the Service forecast an increase in demand for mortgage cases. The Service used its planning tool to determine what action to take. We used our model to run two scenarios to examine the Service's planned corrective actions:

- **Scenario 1**: Increase in the number of mortgage cases arriving per day, from 23 to 39, holding constant the level of resources as at August 2011.
- **Scenario 2**: Increase in the number of mortgage cases arriving per day, from 23 to 39, and with 43 additional adjudicators.

3.4 When simulating Scenario 1, the average number of cases resolved at adjudication fell from 9,673 to 7,651 and the average time a case spends in the system increased from 139 days to 716 days. The system resolved nearly the same amount of cases at Ombudsman stage, but cases spent much longer in the system on average; 807 days compared with 381 days. The increase in the time spent in the system is explained by the rise in the queuing time before the adjudication stages from an average of twelve days to 616 days. This showed that the current level of adjudicators is not enough to process the new volume of incoming cases.

3.5 We then ran the model for Scenario 2, with the increased resources available at the adjudicator stage of the process. The model shows that this has the effect of solving the problems at the adjudicator stage. The number of cases resolved at adjudicator stage, and the average time these cases spend in the system, return to similar levels that were seen before the rise in demand. However, Figure 12 shows how increasing the number of adjudicators does not equally reduce the time spent in the system of cases resolved at later stages. While time to resolve at adjudication shrinks by approximately 82 per cent, for the draft decision and final decision stages time to resolve decreases by 2 per cent and 26 per cent respectively.

3.6 Our analysis shows that the increase in resources has the effect of shifting the backlog to the decision stage in the process (Figure 12). An increase in demand causes a steep increase of time spent by cases queuing in front of the Adjudicator stage, while it shortens the queues at the following stages because cases do not flow through the process (Figure 13). By increasing the number of resources, queues are shifted towards the latter stages of the process, meaning that the overall reduction in time that cases spend in the whole system is limited by the lack of resources at later stages.

Figure 12

Average time spent in the system (days)

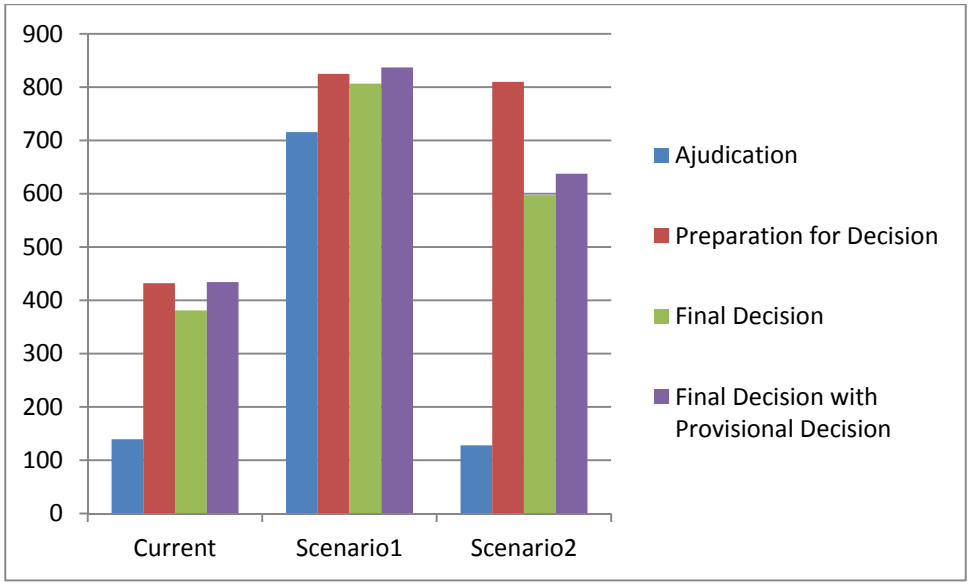

Source: National Audit Office

Figure 13

Queuing times (days)

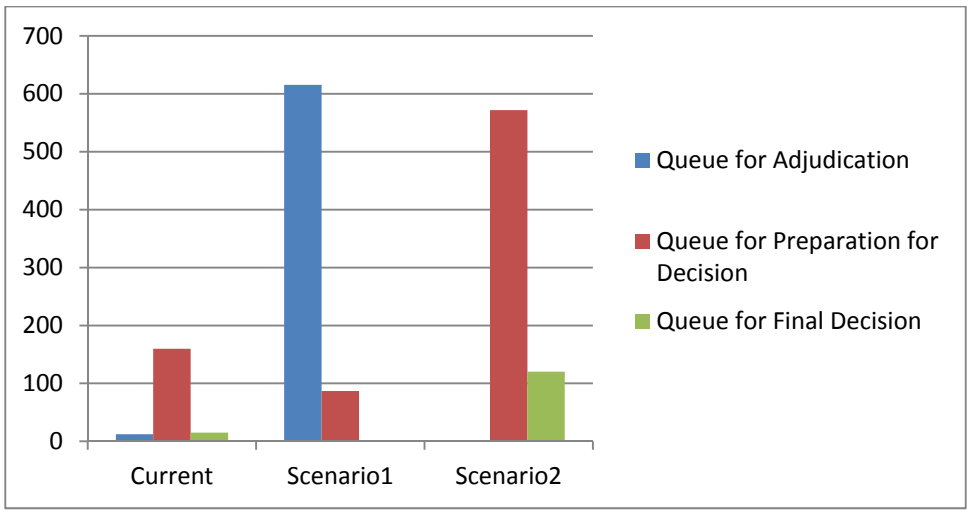

Source: National Audit Office

Part Four

Discussion and overall conclusions

4.1 We developed a model of the Service's end-to-end case handling process using a Discrete-Event Simulation model. Our model simulates how cases flow through the process over time, from case set-up to final resolution at Ombudsman stage. Based on the historical distribution of possible cycle times, the model enabled us to better understand the behaviour of the system over a period of two years to monitor the time spent in the system by single cases, gain insight on where problems might exist and explore modifications to the system.

4.2 We used the model to examine, in particular, mortgage cases. We reproduced the current set up and performance of the process, and then explored impacts of changes in the demand of mortgage cases and the impact of proposed corrective actions that the Service was pursuing. We found that the Service is able to use its planning tool to correctly diagnose the effects of the changes in the demand, and forecast the need for additional resources. In the case of the rise in demand for mortgage cases, the analysis with the Service's tool resulted in the conclusion that an increase in 20 permanent adjudicator staff was needed, and a further 23 contingent staff needed to be recruited for six months to deal with the backlog.

4.3 We used our model to examine the impact of this corrective measure. Our analysis showed that the increase in resources had the effect of shifting the backlog towards the later stages of the process. The model also suggests that the time-frame established by the Service to deal with the backlog created by an increase in demand for mortgage cases may be too short, due to the duration of the overall process. As at February 2011, the Service reports that, in fact, it is ahead of plan in dealing with cases and expects all will be resolved within 18 months.

4.4 By modelling the process in this way, we have been able to highlight some areas in which the Service could look to develop its DC&OP planning tool. Since it uses median cycle times, rather than the distribution of time to process all cases, the DC&OP tool does not fully analyse the impact on queues of the cases that take longer than average to process. This means that it does not currently allow analysis of the impact of volume variations and resource levels on the corporate targets of percentage of cases resolved within both three and eighteen months. These are currently analysed separately, and the Service could therefore seek to integrate its separate analyses of the distribution of elapsed times into the DC&OP tool. The model has the capability to forecast on a longer horizon. However, in respect of mortgage cases, at the time of our fieldwork it was only being used to forecast likely performance against call closure targets six months into the future. The Service could

use the functionality of forecasting further into the future for all products if it judges that the greater analytical effort required does not impede on efficiency.

4.5 Our model also allowed us to examine the sensitivity of cycle times to changes in aspects of the configuration of the system. We found that by altering the maximum number of work-in-progress cases handled by adjudicators, ombudsman support team and ombudsmen at any one time that the average time it takes to process cases through the system can change significantly. This is an area that the Service should analyse further to help it assess the optimal number of cases that can be allocated to staff members simultaneously.

4.6 Currently the DC&OP planning tool can only be used to model the impact of changing the levels of staff resources allocated to dealing with cases, in response to predicted changes in demand. However, the time that a case spends in the process is not only a function of the number of resources available, but also a feature of delays caused by factors that are currently outside the influence of the Service, such as waiting for additional information requests or communicating with consumers through Claims Management Companies.

4.7 Additional resources will reduce queue lengths and queuing times as there will be more staff available to take cases out of the queue, so cases will progress through the system more quickly. However, this will have little effect on the time it takes for a staff member to resolve or progress cases. As the additional resources will increase operating costs, the savings that can be realised from this will only be apparent up to a certain point. Once the queues are empty, an additional resource does not reduce the time cases spend in the system. This point was demonstrated with our model which showed that, when we added the extra resource to deal with the increase in demand of mortgage cases, the utilisation rate of adjudicators was just 78.3% over 2 years, holding an individual's capacity as constant.

4.8 Currently, the overall analysis that the Service carries out across all cases through its DC&OP tool treats each stage in the process as a 'black box'. The timings of the single activities that take place within that stage in the process are not analysed within the DC&OP tool, but are investigated where the Service identifies a problem from the timing of the stage or as cycle times come under examination as part of the continuous improvement programme. A better understanding of how the efficiency of the process can improve, and how the system can cope with increases in demand could be realised through a better understanding of the drivers of delays and flows of information between the different actors in the system. This will allow the Service to target improvement on cycle times and better understand the optimal level of resourcing requirements to meet their performance targets while not increasing costs.