Practice Papers

Optimising retail networks: A case study of locating Camelot's lottery terminals

Received (in revised form): 15th February, 2002

Peter Sleight

founded Target Market Consultancy in 1991 to provide advice on customer targeting and retail location analysis. His previous career encompassed FMCG marketing, advertising (marketing director of Allen Brady & Marsh) and geodemographics — first at CACI, then joint founder of PinPoint Analysis.

Graham Smith

is a senior consultant at Business Geographics Limited, London. His industry experience includes several years providing site location consultancy to companies in retail, media and advertising. He has been heavily involved in work with Camelot during recent years, and was responsible for implementing Camelot's GIS system, 'Optimum'.

Jonathan Walker

graduated from Leeds University's School of Geography and began research on a PhD in retail modelling. Then, in 1998, after working for GMAP (then ULIS), completing an MBA and working for the Rank Group, he set up GeoBusiness Solutions with two colleagues.

Peter Sleight

Target Market Consultancy,
Woodlands, Woodlands
Close, Holmer Green, High
Wycombe, Bucks HP15
6QG, UK.

Tel: +44 (0)1494 712 371;
Fax: +44 (0)1494 714 203;
e-mail: Peter.Sleight@
btinternet.com

Graham Smith

Business Geographics Ltd,
8–10 Dryden Street, Covent
Garden, London WC2E
9NA, UK

Tel: +44 (0)20 7520 5800;
Fax: +44 (0)20 7520 5801;
e-mail: graham@
businessgeographics.co.uk

Jonathan Walker

GeoBusiness Solutions,
The Old Town House, High
Street, Wendover, Bucks
HP22 6EA, UK

Tel: +44 (0) 1296 622 555;
Fax: +44 (0) 1296 622 118;
e-mail: jon@geobusiness.
co.uk

Abstract Camelot has run the National Lottery since its inception in November, 1994. In summer 1998, Camelot decided to investigate whether its existing network of retail outlets within which its 24,500 online terminals were located, was in fact the optimum network — or whether it would be possible to make improvements.

This paper sets down the background to Camelot's retail operation, and how the investigation was conducted. It then goes on to show how a spatial interaction (gravity) modelling approach was devised and implemented, with the aim of providing Camelot with a mechanism to optimise future terminal placement.

Although gravity modelling is not new this particular optimisation model was unusual in several respects. First, against the scale of task. In the more conventional modelling that might be used in (say) grocery retail, a spatial interaction model would typically be used to support decisions relating to a single, or small numbers of possible, predetermined locations. The Camelot model uses a genetic algorithm that allows optimisation across entire regions of the UK at a time, and involves the model working across many thousands of possible retail outlets and millions of possible resulting networks. Secondly, against the usual demand flows (from residential areas to a retail destination) the Camelot model allows for more complex flows (from home, work or shopping trip).

Finally, the paper explains how the model was operationalised within a geographic information system, and implemented for Camelot to run in-house.

BACKGROUND TO CAMELOT

Camelot Group plc won the first contract to operate the UK's National Lottery, with the first draw taking place on 14th November, 1994. The National Lottery quickly became a national institution, and has rarely been out of the news since. The company was built from scratch specifically for the purpose of operating the lottery. It grew from a nucleus of 100 people in May, 1994 to an organisation of close to 500 by the time operations went live. Its seven-year contract ended on 26th January, 2002,¹ at which point it commenced a second seven-year term. The primary remit for Camelot is to maximise revenue for good causes.

Initial selection of retail representation

The initial selection of retailers took place in 1993. The startpoint was population density (from the Census), and a research survey of consumer 'propensity to play' the planned National Lottery. This source provided the demographic drivers of expected lottery playing (for example, it was clear that C1C2D² players would form the key market). The planned online terminal provision was set (at an average of one terminal per 2,000 head of adult population) and was weighted by population profile, on the basis of the research results.

Retailer selection was performed on a postcode sector basis; so each postcode sector was allocated x number of terminals, based on its residential population weighted by demographics. The actual retailers that formed the 'pool' for selection split broadly into multiples (including supermarkets, variety stores, post offices, large confectioner/tobacconist/newsagent (CTN) chains) and independents — largely CTNs, convenience stores and

off-licences. Although the individual outlets of the multiples could offer much larger sales potential than individual independents, it was clearly important to select a 'mix', so that local communities were adequately covered. The independents that had expressed an interest in taking lottery terminals were 'vetted' by a team of ex-Cadbury sales representatives, and agreed criteria were used in the vetting operation. The result of this exercise was that a 'pool' of suitable independents, plus the top 200 outlets from each participating multiple, went into the selection process.

The initial selection, in April 1994, required a total of 24,300 outlets to meet the target retail network agreed in the licence. So a representative of each retail sector was chosen in turn for each postcode sector, until the allocation for that postcode sector was filled. This resulted in 'fair' allocations by retail sector, and an overall geographical coverage. A rollout plan was devised, whose aim was to improve on this initial selection and to fill any gaps that emerged.

IN PRACTICE, IMPROVEMENTS WERE POSSIBLE

Camelot's Retailer Services Division subsequently found that improvements could be made to this initial selection. One key issue was that the startpoint had been residential addresses, whereas some proportion of lottery playing takes place from workplace addresses, or during shopping trips. Camelot subsequently implemented a policy of reallocating some terminals out of dormitory towns, into city centres. Another problem had to do with precise geographical proximity. Some outlets containing terminals could be adjacent to each other on the same high street. Also, there were a few 'black holes' where (in

the ever-changing world of postal geography) complete postal sectors had been missed from the initial selection. Camelot was able to correct for these problems over time by directing allocation of terminals into areas originally missed, using 'on the ground' experience of its sales force.

RULES FOR TERMINAL REPRESENTATION

In discussion with Oflot, the lottery regulator, Camelot agreed criteria for public access to lottery terminals. It has been mentioned that online terminal provision was set at an average of one terminal per 2,000 head of adult population. An additional criterion, volunteered by Camelot, was that at least 95 per cent of the population would live within two miles of a lottery terminal. A third criterion was that there would be at least one terminal within each local authority district; this was the ruling for the first contract, and this was changed in the second contract to at least one terminal per postcode district. Work done during the development of the project described here showed that there were some postcode districts in the Highlands and Islands of Scotland where population density was so low (and retail representation was so poor) that this criterion could not be met. The criterion was subsequently adjusted to read 'within each postcode district of 2,000 or more population'.

SALES EXPERIENCE ALLOWED FOR FINE TUNING

Once the system went live (November, 1994) Camelot was able to take account of actual sales variability by region, and balance terminal representation accordingly. Actual sales experience, and the presence on the ground of the

Camelot sales team, meant that other improvements could be made over time. More precise geographical issues could be addressed, such as adjacency of lottery outlets, and natural barriers to movement could be taken into account. Changes to shopping patterns (eg new shopping centres) and residential patterns (new housing) could be accommodated. So the network evolved over time.

REVIEW IN 1998

In summer, 1998, Camelot Retailer Services called in the first author's company with a brief to:

- review and revalidate the retail network selection methodology
- re-assess the commercial potential for expanding or changing the mix of the network
- review the opportunities for 'starting from scratch' with the most appropriate up-to-date methodology.

Three strands of investigation were followed:

- review data sources within Camelot
- review data/methodologies for 'home *vs.* work *vs.* shop' issue
- review sources for retail location data.

Camelot internal data were examined with two requirements in mind. The first was a consumer-based source of information that could form the basis for a customer demand model; a means for relating lottery playing to geography, using a geodemographic approach. Camelot had a wealth of market research data, covering various aspects of lottery playing, and the retail environment in which it was played. Particularly interesting for the purpose was a tracking study, which interviewed 1,000 consumers each month, identified spend

patterns by demographics, and also asked 'where bought' questions. It should be noted that no name and address information is collected via lottery terminals; so although very detailed data of sales by terminal are known, no transactional consumer data are available.

A second requirement from Camelot's internal data was the database of lottery outlets, including lapsed outlets; and the terminal transaction history from Camelot's data warehouse which would be necessary at the future modelling stage.

POSSIBLE EXTERNAL DATA SOURCES

A number of sources and authorities were consulted about the fairly complex issue of 'flow data' which might be relevant to modelling Camelot's lottery ticket purchases. As stated before, home-based trips to lottery terminals represented the majority of purchase trips; but other purchases were made from work, and in the course of shopping trips. Discussions were held with suppliers of relevant data sets. A particularly useful source of information was the Oxford Institute of Retail Management at Templeton College.

A third data source that would be necessary was the 'universe' of retail outlets that were relevant to Camelot — primarily supermarkets, post offices, multiple CTNs, independent CTNs and convenience stores. There are a number of sources of these datasets, each with strengths and weaknesses. An evaluation of the costs and claimed comprehensiveness of these sources was conducted.

PROPOSED APPROACH

The problem was conceptualised as a demand and supply issue. Consumer

demand for lottery purchases could be modelled using a geodemographic approach, thus relating demand to geographical areas. The physical placement of lottery terminals at specific locations represented supply points. The objective was to match supply to demand in such a fashion that supply of terminals was *optimised*; the simple expedient of putting lottery terminals in *all* available outlets (ie maximisation of supply) would not make economic sense, because the high cost of terminal supply and maintenance would show diminishing returns beyond a certain point.

The recommended solution to the problem was the implementation of a spatial interaction (gravity) model, which would replicate flows from residential areas into lottery outlets; and which could also accommodate work-based and shopping-based flows.

The gravity model should be implemented within a geographic information system (GIS), which would hold the locations and attributes of the retailer universe, area boundaries, population and its basic demographics, and the normal backdrop of geographical features (settlements, 'urban sprawl', waterways and railway lines — that might constitute barriers to movement — road network, etc.). Apart from providing the environment within which the gravity model would function, the adoption of a GIS would also help to solve some of the problems identified by Camelot in their initial outlet selection, *viz*:

- 'black holes' in retail representation
- adjacency of outlets on the same high street
- presence of natural barriers
- identification of 'flows' of shoppers and workers.

USEFUL DATA SETS IDENTIFIED FOR THE PROJECT

As a result of the evaluation exercise, a number of data sets were identified as being potentially useful to the project:

- travel-to-work areas, derived from census data by the National Online Manpower Information System (NOMIS): these areas are reasonably self-contained, so present good study areas for gravity modelling
- CB Hillier Parker National Survey of Local Shopping Patterns: unique source of shopper flow data from residential areas into (a) supermarkets, for main food shop, (b) shopping centres, for non-food shopping
- 'workplace' data, derived from population census data, to show flows of workers from residential areas into their place of work.

In addition to these datasets, candidate sources of retail location data were identified (these were validated at the pilot project stage), as were appropriate boundary files and geocoding datasets (ie means for attaching grid references to addresses or unit postcodes).

CONSULTANCY REPORT TO CAMELOT, AND NEXT ACTION

In summary, it was recommended that a spatial interaction (or gravity) model should be built and calibrated, taking account of flows from residential areas, workplace and shopping trips. A geographic information system should be developed, containing the universe of relevant retail outlets. Census population data should form the basis of a consumer potential model for lottery playing.

The gravity model should be developed initially in a relatively small geographical area (say, the trading area of a large shopping centre), and then rolled

out to cover the country. This development area would also allow alternative data sets to be evaluated.

The model, once developed, would be installed on Camelot's premises, and Camelot personnel would be trained to run it. The GIS environment would also allow Camelot to present visual demonstrations of the optimisation process to senior management, Oflot, etc.

SUPPLIER SELECTION AND DEFINING OBJECTIVES

Proposals were received from a number of potential suppliers; they were evaluated and discussed with Camelot. After discussion and supplementary questioning, Business Geographics Ltd. (BGL) was chosen to develop and implement the model. By agreement, BGL also involved a modeller from GeoBusiness Solutions in model development.

During this discussion stage, the key objectives were presented to the project team. These were to:

- appraise the existing network of UK lottery retailers
- identify the theoretical ideal or 'optimum' network distribution by geographical area
- provide a tool for the ongoing selection of retailers to move towards the optimum network.

So Project Optimum commenced.

CHOOSING THE DEVELOPMENT AREA

The development area was required to cover at least 1 per cent of the UK population and incorporate a wide range of demographic types. Also it was to include at least one major retail centre. This meant that the project team would

be exposed to the full extent of problems that had been identified during the original selection process, and that the model design could address these issues prior to application of the model to the rest of the country.

The Reading travel-to-work area was chosen as the study area for model development. As a part of the development process, an exit survey was designed and implemented to provide information on trip behaviour of lottery ticket purchasers (importantly, the different types of trip in different areas, eg town centre/suburban/rural).

MODELLING CONCEPTS

The model was originally conceived as a variation of the standard spatial interaction model, or 'gravity' model, a mathematical approach which is well established as one of the more sophisticated tools available to retail location planners.

Spatial interaction models provide a computer-based, mathematical representation of the following elements for any market:

- demand (size and type of population)
- supply (eg location of retail outlets)
- the 'attractiveness' of retail outlets
- the effects of distance or travel time (accessibility to retail outlets).

These model inputs are associated with geographical coordinates and thus effectively spread over space as 'layers'. By calibrating from known trip patterns, the model replicates the interaction between demand and supply as flows of people or, in this case, spend on lottery tickets.

By aggregating the results of such a model, the overall effectiveness of a retail network can be determined in terms of the movements required for demand to

be satisfied, and the network can be altered within the model until some idealised situation is reached; that is, the 'optimum network'.

Although spatial interaction modelling is not new, this particular model is exceptional for a number of reasons. In the more conventional modelling that might be used in, say, grocery retail, a spatial interaction model would typically be used to support decisions relating to single, or small numbers of possible, predetermined locations. The Optimum model would need to take into account many thousands of possible retail outlets. While conventional models are used to test the impact of single prescribed scenarios (eg opening a new outlet, closure of an outlet), the Optimum model would be used to determine the optimum distribution of the entire network, involving millions of possible combinations. The model is also unusual in terms of the flows that it represents. Conventional models are typically concerned only with flows from residential areas to a retail destination. The Optimum model has been designed to allow for the more complex flows associated with work and shopping trips in matching the location of demand.

CAMELOT SELECTION DISTRICTS

Another original feature relates to the units of area used in the modelling process. Given the relatively short trips that are typical (measured in hundreds of yards rather than miles), postcode sectors would be much too large. (Postcode sectors average about 2,500 households, there are some 9,000 in the UK, and in practice they vary enormously in geographical area.) Census Enumeration Districts (EDs), which average about 175 households each, were of a more suitable size, but in dense urban areas they are

Figure 1

too small and would result in unrealistic model complexity. So a new geography was developed for the purpose — Camelot Selection Districts (CSDs), which were aggregations of EDs; thus, an intermediate geography between EDs and postcode sectors. To give an indication of the relative sizes of each, the Reading development area covered 58 postcode sectors, or 687 EDs. After aggregation, a total of 238 CSDs provided a suitable middle ground.

Four types of CSD were defined, each with distinct characteristics, either as origin zones or destination zones, or both:

- type 1: town centre
- type 2: suburban
- type 3: rural
- type 4: supermarket (pseudo-CSD, not really a district but a distinct destination).

Figure 1 shows CSDs to the north-west

of Reading, together with population-weighted postcode centroids — the points at which demand within each CSD would be represented in the model.

PLAYER DEMAND: THE CONSUMER POTENTIAL MODEL

The National Lottery network is in two strands. The OLT network, or ‘online’ network, is a network of terminals operated by retailers selling lottery tickets for the Saturday and Wednesday draws. The online game accounts for around 90 per cent of all lottery sales, and is therefore the primary concern. In addition, the Game Validation Terminal network provides sales points for ‘Instants’ scratchcards. Instants can be purchased from OLT retailers, and from a further 10,000 retailers not selling online tickets. Although a separate model was developed for Instants, which represent more of an ‘impulse’ market,

this paper is concerned solely with the work conducted on the online model.

Camelot's telephone tracking survey provided information on frequency of spend and average spend per purchase for each respondent. These values in turn would be used with demographic data for the test area to provide an estimate of the spend within each CSD.

Since social grade (AB, C1, C2, DE) had been recorded for all respondents in the tracking survey (23,807 records), this could be linked to 1991 census data and converted to a model based on social grade. The census provides data on the number of people in each of 17 'socio-economic groups' (SEG) for each Census ED. These SEG groups have been translated to equivalent social grade categories.

This option would involve averaging all values for respondents within the same social grade. The assumption would be that the average spend for respondents within any one social grade would be the expected spend for all people in the test area that were classified within that social grade.

REPRESENTING TRIP PATTERNS IN THE MODEL

In constructing a spatial interaction model, observed trip patterns must be obtained for model calibration. The term 'trip pattern' refers to the nature of consumer movement: what effect does distance have on the propensity to travel in order to make a purchase at a particular location? How far do people travel to buy lottery tickets? During calibration, model parameters are adjusted until the predicted flows within the model most closely replicate actual trips. In other words, when the average distance of flows between CSDs most closely reflect actual average distances travelled. As no such data were available,

an exit survey was carried out at lottery outlets in the Reading area.

Trip distances were calculated for each respondent between the point of origin and the retailer at which the lottery purchase was made. The 'point of origin' was defined as the location from which the respondent set out to purchase, be it home, a friend's home, work, a particular shopping location, and so on.

In order to obtain input values for model calibration, results from the exit survey were divided into eight categories, accounting for different trip types. These would produce eight separate 'beta' values to represent distance decay effects in the model. The categories were classified according to the type of area from which each trip originated (urban or rural) and the type of destination in which the lottery purchase was made (town centre, suburban, rural or supermarket).

Spend-weighted drivetimes (for town centre and supermarket destinations) and spend-weighted road distances (for suburban or rural destinations) were derived for each of the eight categories and used as the trip cost for model calibration.

A further highly valuable outcome of the exit survey is that it provided information on the split between lottery spend which could be attributed to trips from home and that which could be attributed to trips from work. Using this information it was possible to produce composite demand values for all CSDs, based on the number of people living in a particular area (predicted demand weighted by social grade) and also the number of people working in that area.

SHOPPER FLOWS

Shopping-based demand provided the most difficult challenge for representation in the model. While home-based and

workplace-based demand can be attributable to particular points of origin and tend to involve short trips, shopping is much more transient in nature.

The main element to shopping-based demand in the 'Optimum' model is shopping in town centres (also out-of-town shopping centres). These are CSD 'Type 1' locations. The demand for lottery tickets associated with town centres is treated in a wholly different manner to home/work-based demand within the model.

As a point of note, supermarket shopping was treated as part of the home/work-based demand within the model. Supermarket destinations were given greater pulling power within the model to reflect the greater distances that people will travel to supermarkets than to other retailers (eg CTNs, convenience stores).

Shopping at small local centres was also accommodated by the home/work-based element of the model, as a part of the general movement within the model.

The source of data used for analysis of town centre shopper flows was the 1998 'National Survey of Local Shopping Patterns' from surveyors CB Hillier Parker.³ This is a national data set based on the shopping habits of approximately 1.2m respondents. Survey results have been modelled at postcode sector level to provide estimates of the number of people in each postcode sector that most regularly visit one of around 2,500 town centres for non-food shopping. These were aggregated to produce values for the 'shopping population' of each centre.

In order to account for the importance of shopping destinations in the spatial interaction model, terminal numbers were pre-allocated — that is, determined as suitable for a given number of terminals, before running the optimisation routine. This was by a

detailed analysis of the sales that could be supported by different sizes and types of centre.

MODEL STRUCTURE AND OPTIMISATION ROUTINE

The spatial interaction model was developed as a Visual Basic application. Demand and supply were represented at the CSD level as tables in a separate database. This database also contained other tables essential to running of the model, including a matrix of trip costs between CSD, and various tables for model output.

An optimising genetic algorithm was developed within the model to score successive randomly generated networks and mutate towards the network that provided the best balance between supply and demand across geographical areas. The methodology for scoring each network was based on the innovative concept of 'effective supply' of terminals. By backpropagating modelled sales at each supply point to source, levels of demand and effective supply could be compared for each CSD within the general formula.

A terminal located in any one destination zone (CSD) is not only used by people living, working or shopping in that zone, but also by people in neighbouring zones. Because of this, the 'effective supply' is the number of terminals actually serving an individual zone. By way of example, if a terminal in one zone pulls custom in equal volumes from four CSDs then each CSD has an effective supply of 0.25 terminals.

Using the effective supply principle, all terminals in any given network configuration can be reallocated to the CSDs which they effectively serve. The total of the effective supply across all CSDs will always equal the actual number of terminals. Since the number

of CSDs in the UK is greater than the number of available terminals, and since some CSDs, notably town centres, clearly justify multiple terminals, many other CSDs will not have a terminal. The reallocation process ensures that all CSDs have both a demand term (expected spend) and a supply term. Supply then becomes a continuous layer in the model, as with demand.

Having established the demand and supply for each CSD, the optimising model seeks to balance the two values. The optimal solution can be defined in modelling terms as that which minimises the variance between the average ratio of supply to demand and the actual ratio in each zone. An optimal network will be one where the ratio of demand to supply in each CSD is as close to the average ratio as possible. Such a network would mean that each CSD is equally well served by lottery terminals, regardless of whether terminals are physically located in a particular zone.

The method adopted takes a starting network containing a predetermined number of terminals. The starting network is unimportant, but one that roughly reflects the distribution of demand is sensible. It will also be constrained by the location and number of supermarket and town centre terminals, which were pre-allocated by separate modelling processes, the detail of which is not included here, and by the number of suitable retailers within each CSD (the retail 'universe').

The starting network is run through the model and, based on the predicted trip patterns, the effective terminal supply for each CSD is calculated. A score is calculated for the network, by taking the sum of all the absolute variations from the average ratio between effective terminals and demand. The lower the score, the better the network.

After the network has been scored,

random changes are made to the network. The model is then run again and scored. This process is repeated several times within a single iteration of the optimisation routine, and each score is automatically logged. After each iteration, the network with the lowest score is taken as the new starting point. Random changes are made to this network and so on. The model can be regarded as based on evolutionary theory because, as in true evolution, random changes that produce poor (or poorer) networks are lost whereas random changes which improve the performance of the network become incorporated in future generations of network.

The graph in Figure 2 shows how the 'best network score' steadily improves with many iterations of the optimising routine.

DEVELOPMENT AREA RESULTS AND NATIONAL OPTIMISATION

The results for the Reading area showed the idealised network density distribution to be somewhat different to the existing network, but not too dissimilar. A thorough comparison of the existing and 'optimum' network was presented to Camelot, and proposals were also discussed regarding the modifications that would need to be made to translate to a national model.

The model was then rolled out to the remainder of the UK; it would have been too unwieldy to run the model for the whole of the UK all at one time. It was therefore useful to divide the country into regions, with the model run separately for each region. Unfortunately, government standard regions are not composed of whole travel-to-work areas, but instead are made up of counties. In order to avoid problems of travel-to-work areas being split between regions, model regions were produced

Figure 2 Improvement of network score by iterative optimisation

that were as close as possible to government standard regions, but which were composed of whole travel-to-work areas.

With 25,000 terminals available for the network, this was split according to demand across the regions. Further development of the basic development model was necessary before national optimisation (particularly to accommodate London), but once work was complete, the model was run on each region in turn, and the results of optimisation collated to produce the final idealised network distribution for the UK, defined across 48,000 CSDs.

DELIVERING THE MODEL RESULTS: DEVELOPING THE GIS SYSTEM

Having established the theoretical optimum network, the project team were required to develop a desktop GIS system for Camelot to access the results of the model, and provide a tool for

network planning. The system is a heavily customised adaptation of an existing 'off-the-shelf' mapping software. The GIS database holds the 'retail universe': a unique and constantly updated database of all convenience and other retailers in the UK, with existing lottery retailers flagged. 'Retailer selection reports' are generated using these data to provide sales executives with information relating to retailers in any specified area, as well as a map showing areas of oversupply and undersupply of lottery retailers based on the Optimum model.

BENEFITS TO CAMELOT

For the 150 field sales executives who visit and make selections of new lottery retailers, while local knowledge is still vital, the provision of reports from Optimum has dramatically improved the efficiency of the selection process.

At Camelot headquarters, the system has also improved the organisation's

strategic network planning capabilities. Users of the system, itself known as Optimum, can now easily produce statistics, reports and maps for any single area, or for the UK at any level of geography, relating to demographics, demand, terminal numbers (existing and optimum) and lottery sales, as well as market statistics combining lottery data with demographics (eg sales per player). It also allows monitoring of other aspects of the network from a geographical perspective, including outlet proximity (to other features), provision for rural communities, split of outlets between retail types, and so on.

The primary role of the retailer services division is to support existing lottery retailers and ensure the smooth operation of the lottery network as a whole. As operators of the nation's lottery, however, Camelot also provide a service to the general public. Indeed, over 60 per cent of the adult population (approximately 30m people) still play the lottery regularly — one of the highest participation rates anywhere in the world. The definition of 'customer' therefore includes catering for enquiries and requests from players, local councillors, MPs, and others. Being able to provide accurate information, and quickly, is vital to customer confidence, and Optimum has greatly enhanced this ability.

As an example, The Rural Shops Alliance recently approached Camelot to understand what was being done to support rural communities and independent retailers. A member of the selection team was able to demonstrate Optimum's functionality, and the work that was being done to ensure that rural retailers continue to be an essential element of the lottery network. Feedback indicated that they were 'deeply impressed'.

Crucially, Optimum now offers

Camelot a greater degree of objectivity in retailer selection, and the importance of this cannot be understated. Although sales executives ultimately make the decisions, over 4,000 telephone calls and 200 letters per month are received from retailers (ie not including calls from the general public), and these are mainly requests for a terminal. Camelot can now provide a much more transparent account of the selection process, as well as the very limited availability of terminals. In efficiency terms, the time taken to select new retailers was halved.

Other tangible benefits that have arisen since Project Optimum was implemented include:

- identification of poorly served communities: the accessibility study carried out by the project team revealed that 96.2 per cent of UK households are within two miles of their nearest online lottery terminal (84 per cent within half a mile). Yet looking closer it was revealed that a number of small rural communities were not that well served, each with clusters of households more than five miles from a terminal. Priority was given to providing these communities with local access to the game
- ringfencing of 'community outlets': the project has enabled Camelot to segment the lottery network, and this includes the identification of 'community outlets' — essentially covering retailers that provide the online lottery to households that would otherwise be remote from the network. These outlets are now given special consideration
- improvements to the network: in relation to the network itself, and movement towards the 'optimum' distribution, it has to be said that improvements to the network are very difficult to identify in the real

world, given the density of outlets. The lottery network, while not ubiquitous, is the largest retail network in the country⁴ (with sales of around £4.5bn per year, the online lottery game represents the biggest single consumer brand in the UK). One simple measure is the distance of households from their nearest terminal. While this does not allow for a consideration of transient demand (shopping, workplace etc, which is accounted for in the model), it does provide an unarguable statistic. The network is already very strong, and it is not easy to value the benefits of further improvement, although to offer some perspective of the rewards on offer, an increase in annual sales of just 0.1 per cent would mean an extra £35m revenue during the next licence period, based on current sales.

CONCLUSION

The Optimum model can only provide an idealised representation of the optimum lottery network, and ultimately there is no substitute for local ground knowledge when making individual site location decisions. Site factors, such as a

transport hub location, or the seasonal effects of tourism, are recognised as being beyond the scope of the model (at least at the moment). At the strategic level, however, the Optimum model provides Camelot with a much more accurate representation of the ideal lottery outlet density than could be realised with the original distribution policy.

In conclusion, the authors have shown how the brief to optimise Camelot's lottery terminal network was approached, and how the solution was devised and implemented. They are not aware of other network optimisation projects of this scale; indeed, the judges of the 2001 Information Management Awards agreed, and gave Project Optimum a Silver Award in the GIS category.

References

- 1 Camelot's contract was due to end in October, 2001, but the delay occasioned by the change of Regulator extended the first contract by three months.
- 2 This refers to NRS Social Grade, ABC1C2DE.
- 3 CB Hillier Parker, 77 Grosvenor Street, London W1K 3JT.
- 4 The second largest network in the UK is believed to be the Post Office network, with 18,000+ outlets at time of writing.