Global optimisation as an aid to resolving consumer price obfuscation: mobile phone tariffs

Les Hatton
CIS, University of Kingston*

December 16, 2004

Abstract

As consumer society becomes more and more complex, finding the most competitive price for a product or a service itself becomes correspondingly complex. This is particularly true of services and in some consumer areas, the complexity and rate of change are such that no human could give consistently good advice. In this paper, some of the relevant mathematical strategies will be reviewed and a case history involving the building of a web resource with a computational engine to find the cheapest mobile phone tariff (cell phone plan in the US) within supplied constraints will be described. Data from both the US and the UK on average savings across around a thousand users will be shown to give an idea of a) how much can be achieved using sophisticated techniques and b) the vulnerability of the public to excessive price obfuscation.

$Date: 2004/12/16 19:00:16 $

1 Overview

The author's interest in this subject was first piqued by a visit to a high-street mobile phone [1] shop claiming to be independent of any of the tariff suppliers, to change tariffs for four mobile phones after experiencing unease about the costs of running the phones on their existing tariff, a business deal treating all four phones together. The salesperson went through some of the many possible combinations before recommending relatively emphatically one particular newly announced tariff which was "a really good deal".

The author wary of such unsubstantiated claims amidst the noticeable drop in basic numeracy of the past two decades, decided to take all the information home and calculate it based on the last two monthly itemised billings. The end product of several hours of patient calculation was that the salesperson was

*L.Hatton@kingston.ac.uk, lesh@oakcomp.co.uk

[1] For US readers, a UK "mobile phone tariff" is synonymous with "cell phone plan". There are also strong similarities in the way they are priced but with some important differences. From here onwards in this paper, the phrase "mobile phone tariff" will be used interchangeably with "cell phone plan". Many other parts of the world use similar nomenclature and very similar pricing strategies.

completely wrong and that one other tariff the author tried appeared significantly better, (by a factor of around 40% !). Spurred on by this, the author decided to look at all the other tariffs before the realisation that there were then over 80, each with many sub-options rendering the whole thing simply infeasible by hand. At this point, the seeds of this paper began to germinate.

Deliberate and accidental obfuscation It is a moot point just how much of the obfuscation in an area like mobile phone tariffs is deliberate and how much is accidental. It is certainly to the benefit of the suppliers to obfuscate on the grounds that if there were ever a clear, unambiguous and easily worked out best deal, most people would take it given the global outreach of the web. Deliberate obfuscation helps to muddy the water to avoid any decision being obvious and enabling a consumer to choose even a very over-priced alternative through inability to see through the layers of complexity. This might be considered somewhat cynical so it should also be noted that when things are extremely complex, accidental obfuscation is also likely when the supplier doesn't really understand the full complexity either. (There were a number of examples in the case history below where a 'special deal' was more expensive than the alternative it was supposed to replace. The nature of these looked accidental but may of course not have been.)

Rate of change Another significant factor in obfuscation is the rate of change. In mobile phone tariffs, it eventually proved necessary to update the tariffs or at least check they were up to date every month and it was a rare month when something did not need changing. The pace of change also accelerated near festive times such as Christmas where an almost weekly check had to be made to stay current. This accelerated rate of change coupled with the fundamental complexity caused by the presence of seven or so suppliers (in the UK) locked in deadly competition each offering bewildering numbers of options and add-ons means essentially that *only* a computational server using sophisticated mathematics was able to track them all.

1.1 Global optimisation strategies

In essence, there are two possibilities for finding a global minimum in an (N+1)-dimensional function $F(x_1, x_2, ..., x_N, t)$ at a particular time t:-

- Exhaustive searching. Here, every possibility is explored and the smallest one found.
- Non-exhaustive searching. Here, usually for reasons of time constraints, a systematically chosen subset of the domain of F is searched in the hope that the global minimum or at least one close to it can be found. In the case of consumer product pricing, this is complicated by the fact that the optimisation surface is discontinuous due to the arbitrariness with which tariffs are priced.

Which of these two approaches is used depends entirely on how much computation can be afforded within the available time. First of all, the time dimension is fixed as the rate of change of tariffs is of the order of weeks rather than minutes, (although with an eye to the future, this may not always be the case). Early versions of this project were written in TclTk, [1]. An exhaustive search of the complete domain of F() for UK mobile phone tariffs could take as much as a minute even on a relatively high-end server, which is far too long for an interactive service. In the early days of the project, the domain was then searched using a *simulated annealing* approach, [2]. In simulated annealing, domain values $(x_1, x_2, ..., x_N)$ are chosen randomly and the function value F() calculated. If F() is bigger than the current minimum, it is rejected. If F is smaller than the current minimum, it might be rejected based on the value of $e^{-\phi T}$, where $\phi > 0$ is a constant and T is the 'temperature', a value which gradually diminishes during the minimisation run. The idea here is that there is quite a high chance that a value will be rejected at the beginning of the run even if it is lower than the current minimum. This is to prevent the algorithm getting stuck in a local minimum whilst missing a better one somewhere else. The probability of 'jumping out of a local hole' gradually diminishes as the algorithm progresses.

Simulated annealing is very effective when the number of possibilities to be explored dwarves those for which there is time to try and initially this met with considerable success. However, just as an experiment, the computational server was re-written in C and the resulting increase in speed was so dramatic that an exhaustive search could be carried out in about half a second which was entirely sufficient for an interactive web request. As a result, global optimisation strategies were not considered further but may well become necessary as obfuscation increases, (which it did inexorably throughout the duration of this study). For a good general discussion of a number of alternatives to annealing, see [3].

Degree of uncertainty The degree of robustness of the solution is an important part of any optimisation strategy with noisy and inconsistent data. The accuracy to which a user can estimate their own phone use can be subject to quite wide variations. In addition, some of the necessary parameters, for example the percentage of all calls made between different network operators which in some regimes like the UK attract different rates, is usually unknown. The network operators do not reveal this information on their itemised billing. To cater for this, a degree of uncertainty was introduced into the user interface in plain English terms by asking for the degree of confidence between "Very confident", "Reasonably confident" and "Not very confident". These were used to generate corresponding degrees of randomness in the searching engine described below.

2 The minimisation criteria and problem parametrisation

The minimisation goal of finding the cheapest mobile phone tariff for a given level of confidence in the input data is quite simple but understanding the problem parametrisation is considerably more complex.

Complexities in mobile and cell phone pricing Complexities are numerous. They include but are not limited to:-

- There are typically between 5 and 10 suppliers of tariffs or plans in most countries. Most suppliers have multiple plans and in the UK this led to initially 60 or so which grew to over 80 in the years of the study. In the US, there were about 120 at the end of the study.
- Each plan has numerous and frequently different options, add-ons and volume related discounts.
- The plans change about every 3 months with increased change near seasonal periods like Christmas where the imagination of tariff architects knows no bounds. Some changes are pathological with entire tariffs being scrapped and replaced with new-look ones. Sometimes changes are subtle with details remaining the same but prices dropping by some percentage.
- The treatment of text or SMS messages, a massive growth area, is very different between suppliers. Some suppliers offer bundles of messages which must all be used, others offer graded bundles with step function pricing. These have on occasions been so complicated that the suppliers have accidentally made mistakes so that a 'special deal' is more expensive than the option it is supposed to be replacing.
- Multiple phone tariffs allow the users to share a pool of speaking or text message time.
- Some tariffs treat all network calls as being equivalent whilst others use price differentiation.
- Some tariffs offer a flat rate for the first N minutes of use and then discounted rates for all or some calls after.
- Special services such as sports scores, internet use are priced somewhat differently to normal speaking or text calls but not always.
- Some tariffs are monthly based with a minimum period whilst others are pay per use.
- Some tariffs do price differentiation on distance between the two subscribers, (for example India and the USA). Others like the UK do not.
- New features with sometimes completely differing pricing structures to the above, for example for e-mails.
- Some countries (like the USA) price differentiate between the cost of sending a text (SMS) message and the cost of receiving one.
- Roaming charges particularly across country borders vary dramatically and some tariffs attempt to cater specially for this.
- The time period deemed 'peak time' is frequently different between suppliers and some suppliers further differentiate between evening and weekends.

- There is a rapidly growing 'middle-man' business area where companies buy time from the suppliers and sell on solutions for multiple phones to corporations.

This list is not exhaustive but should be sufficient to convince the reader that the complexity is already staggering and the chance of a salesperson giving reasonably correct advice is basically non-existent. Not only that but the entire infrastructure is maintained by artificial complexity because if there was a simple and obvious comparison, most people would simply adopt it.

3 Architecture of a web service

The basic idea was to set up a simple web server which solicited enough target information from the user for a given level of use of a mobile phone, either actual (an existing user) or planned (a new user) and some idea of how uncertain these estimates were. In the end, a web service using PHP was constructed which acted as a client, driving a server written in C for efficiency which actually performed the calculations before passing them back to the PHP client for display.

Questions for the users Users were differentiated into two types:-

- Existing users: people who already had a phone and had a reasonable idea of how they used it. Such users would be solicited for their average monthly bill.
- New users: people without a phone who had only a vague idea of how they might use it and had no real idea of actual costs. Such users were not of course solicited for their average monthly bill.

In addition, both kinds of users were solicited for answers to the following questions:-

- Average number of minutes used per month for calls
- Average number of text (SMS) messages sent per month.
- Percentage of all calls made at peak times
- Percentage of all calls made to normal landlines
- Percentage of all calls made to the same network supplier
- Percentage of all calls made to a different network supplier
- Percentage of all calls made whilst roaming
- An estimate of the confidence in the above answers

Note that percentages were solicited as "None", "Some", "Half", "Most" and "All" to cater for any numerical phobia. Confidence was solicited as previously described.

It was found that 22 variables was sufficient to give good robust parametrisation of the above complexity. These variables and their values were encoded as text files, one for each tariff with options for that tariff included within the text file. This allowed the tariffs to be updated easily to keep up with the rapid change in most cases but twice in the one year period, tariff changes forced upgrades in the way the optimisation engine actually handled the data.

3.1 Data handling

Before submitting to the computational engine, the data had to be edited. Some web-users (about 2% in the sample here) delight in entering crazy values for parameters. Obviously stupid ones were simply removed. Furthermore, it was decided arbitrarily that anybody spending more than about 150 pounds (around 300 dollars) per month should probably think more about the danger of frying their brains than of saving money, so larger amounts were simply ignored. This removed another 1% or so.

3.2 Security

PHP was used to drive the computational engine. Several denial of service attacks were attempted even though this was a free service. The service was fully check-pointed with mailers to show where and what users were doing so as the DoS attacks continued, the profiling was used to design counter measures. A combination of checking DNS values, putting up random words and asking users to input 2 randomly situated letters which were then compared with the internal PHP session management values and incorporating exponential queueing in the server response proved sufficient to keep these under control for the duration of the service.

Malicious entries were occasionally entered into PHP fields to attempt to gain control of the server. These were protected against in the usual ways by checking all inputs for malicious content, attempts to execute arbitrary code on the server and attempts to include spoofed PHP scripts.

4 Results

The savings are extraordinary revealing the degree to which the public is vulnerable to price obfuscation. Approximately 1000 unique users took advantage of the service over a 12 month period. Most of these involved the UK computational server but some early data from the US server is also shown for comparison.

4.1 UK mobile phone tariffs

The distribution of savings for each of the four quarters the service was run for UK phone tariffs is shown in Figures 1,2,3 and 4. Taken across *all* users, something like 80% could have saved around 300 pounds (around 600 dollars) *each* per year. Furthermore, it can be seen that the savings increase as time goes by and the phone charging schemes become ever more complicated.

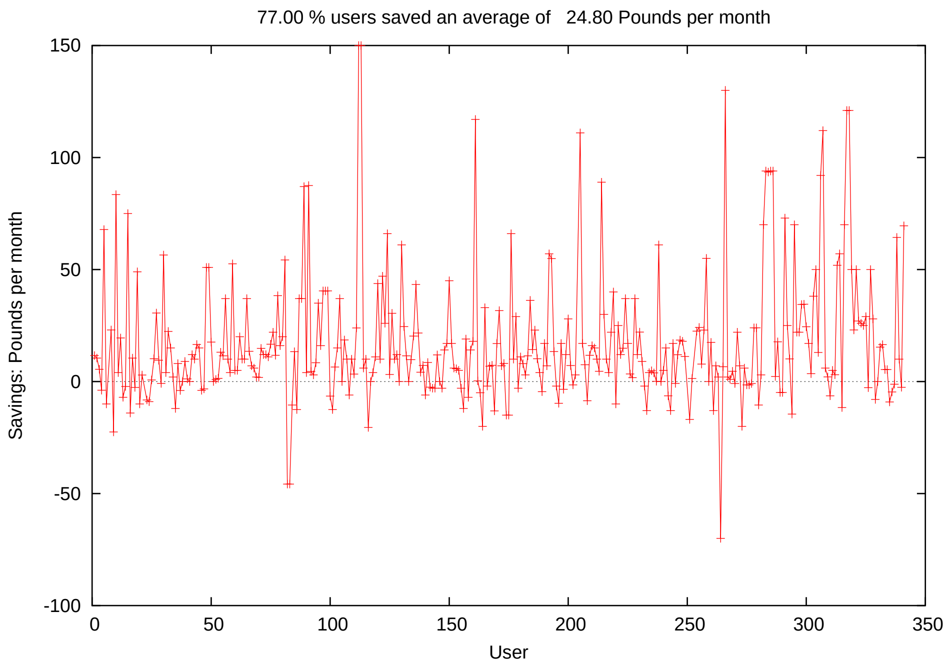

Figure 1: UK server: Savings per user, pounds per month, fourth quarter 2003

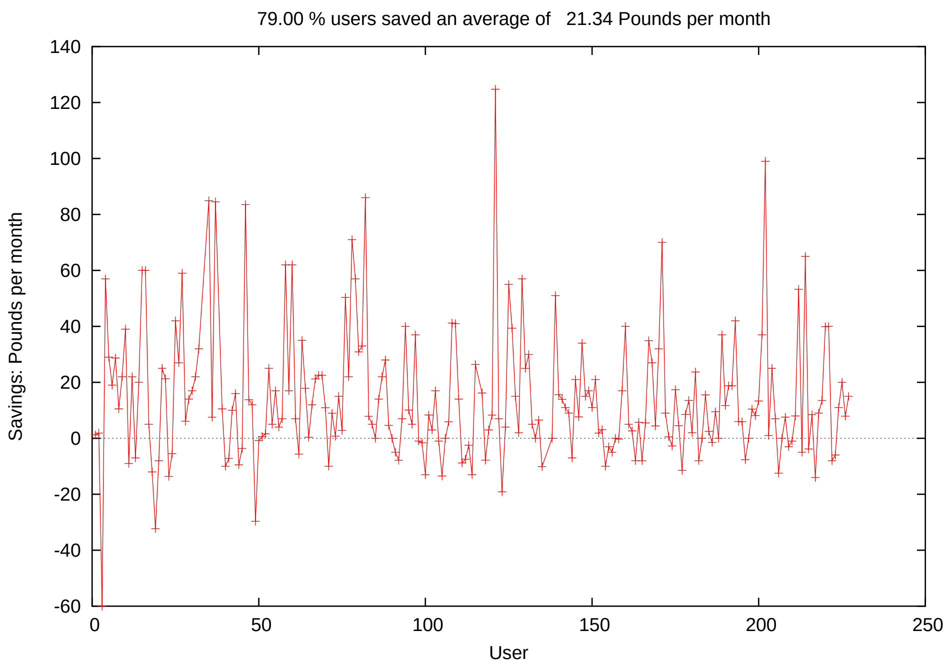

Figure 2: UK server: Savings per user, pounds per month, first quarter 2004

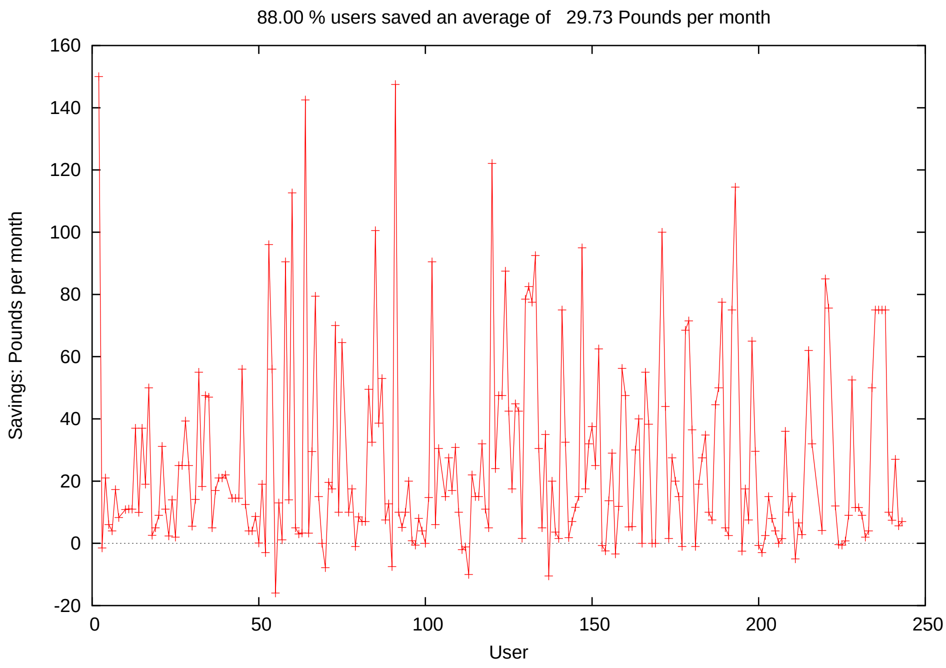

Figure 3: UK server: Savings per user, pounds per month, second quarter 2004

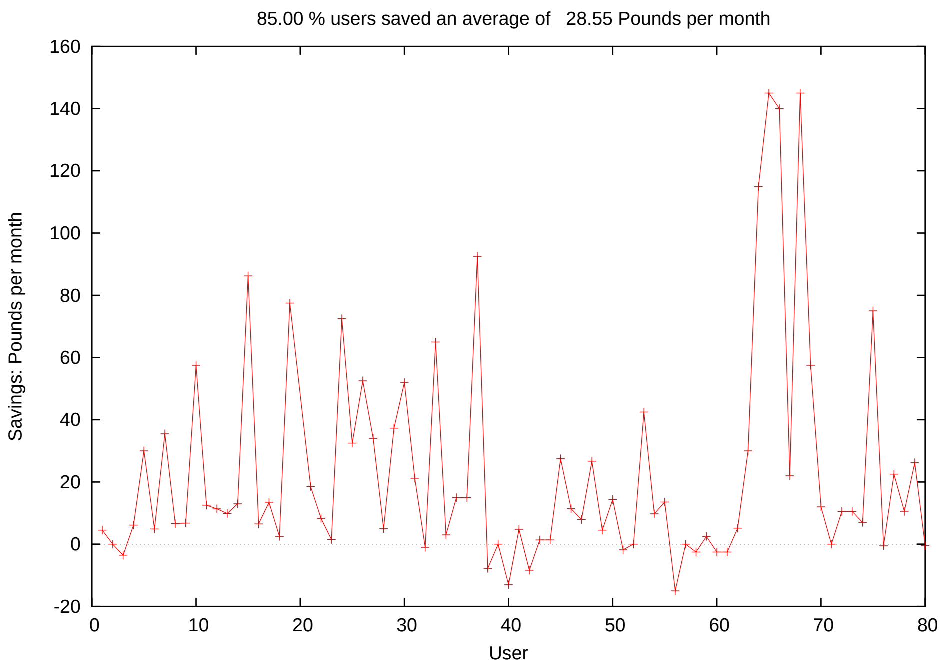

Figure 4: UK server: Savings per user, pounds per month, third quarter 2004

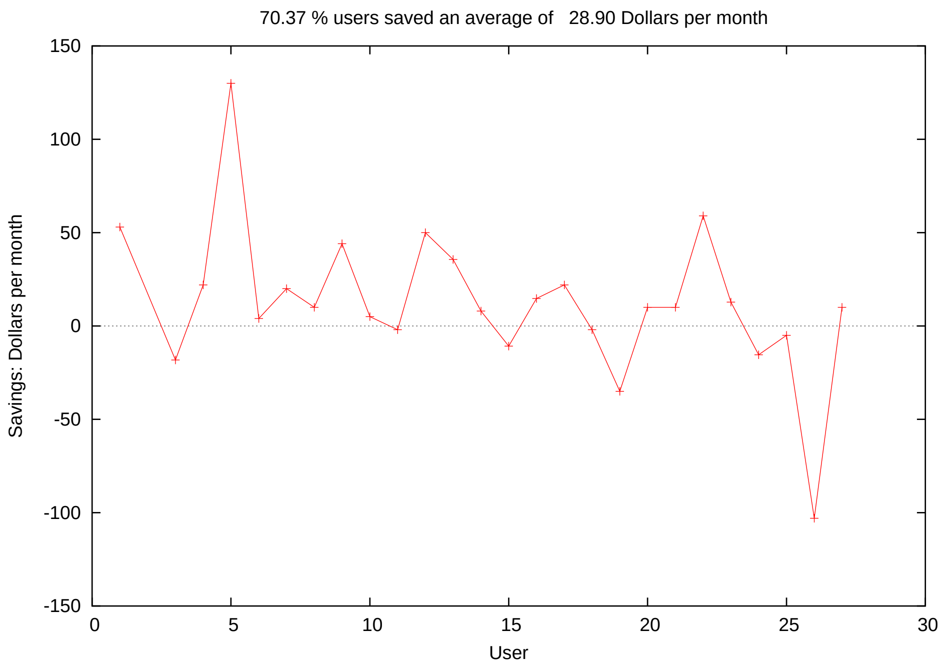

Figure 5: US server: Savings per user, dollars per month, fourth quarter 2003

4.2 US cell phone plans

The picture in the US appears very similar as can be seen by looking at Figure 5. The US site was only run for a short period as it proved too difficult to keep up with the vagaries of phone charging in two countries with the time the author had available, particularly as the Christmas period encroached. However, the data collected suggested a very similar picture in the US with around 70% saving an average of around 360 dollars per year each.

Personal use To make sure these were not hypothetical, the author changed tariffs of four phones according to the predictions made by the computational server. Average savings since changing have been 45% on the original tariffs.

5 Conclusion

This paper introduces the notion of using computational anti-obfuscation servers. As the consumer world becomes ever more competitive, options for the same service increase and complexity increases even more dramatically. The effect on the average consumer is that they finish up paying much more than they need simply because of this increased obfuscation. Furthermore, as can be seen in Table 1, the effect seems to be getting worse.

Quarter	Average savings per month (pounds)
4Q03	24.80
1Q04	21.34
2Q04	29.73
3Q04	28.55

Table 1: Average savings per month per user as a function of time

This paper demonstrates that it is entirely feasible to track this complexity and to control it using standard mathematical techniques deployed in novel ways as servers which anyone can access. The techniques are equally relevant to other areas in which consumers are vulnerable to overly-complex charging such as in financial services.

References

[1] Tcl Tk http://www.tcltk.org/

[2] Kirkpatrick S, (1984) *Journal of Statistical Physics*, vol 34, pp 975-986

[3] Press W.H., Teukolsky S.A., Vetterling W.T., Flannery B.P. (1999) *Numerical Recipes in C*, Cambridge University Press ISBN 0-521-43108-5