



# British Crime Survey: Methods Review 2009 *Final Report*

Authors: Sarah Tipping, David Hussey, Martin Wood and Jon Hales

October 2010

Prepared for the Home Office



# Contents

<b>Executive summary .....</b>	<b>5</b>
<b>1 Introduction .....</b>	<b>9</b>
1.1 Previous methodological studies related to the BCS.....	9
1.2 The BCS Methods Review, 2009-10 .....	11
Sampling .....	11
Response rate .....	11
Reporting period and reference period .....	11
1.3 Approach to conducting the methodological research .....	11
<b>2 Evaluating changes in the BCS sample design.....</b>	<b>13</b>
2.1 Changes to the BCS sample design.....	13
Disproportionate sampling .....	13
Clustering.....	14
Sample size.....	14
2.2 Estimating the effects of changes in sample design .....	16
The impact of disproportionate selection probabilities .....	16
The impact of changes in the degree of clustering .....	19
2.3 Summary of findings.....	21
<b>3 Response rate for young people in ‘high-crime’ areas.....</b>	<b>23</b>
3.1 Background and approach .....	23
Background .....	23
Approach.....	23
3.2 Broad patterns of victimisation on the BCS .....	24
Overall levels of crime.....	24
Young people and crime .....	25
Prevalence of violence by age within sex .....	26
3.3 Victimization by age and area type.....	27
Prevalence of violence by age and area type .....	27
Incidence of violence by age and area type.....	29
Assumptions.....	31
3.4 Response rates by age.....	31
Response rates for young people .....	33
Response rates for young people in deprived/inner city areas .....	35
3.5 Response rates by area type (and interviewer observed variables) .....	36
Categories with low response rates .....	37
Comparison with other surveys.....	38
Modelling non-response.....	38
3.6 Effectiveness of calibration weighting.....	39
3.7 Modelling victimisation (violence).....	39
Characteristics associated with increased risk of violent victimisation.....	41
3.8 ‘What if’ analysis.....	42
3.9 Summary and conclusions .....	43
Summary.....	43
Conclusions.....	43
<b>4 The period covered in annual reports .....</b>	<b>44</b>
4.1 Background and task.....	44
4.2 Calculating annual incidence rates.....	47
A rate per month approach .....	47
Assigning incidents to months.....	47
Calculating annual estimates from the monthly rates .....	49
Weighting issues .....	49
Significance testing .....	50
4.3 How do the two measures react to changes in crime rates in a period? .....	50
4.4 The crime trends.....	54
Comparing the trends.....	54
4.5 Conclusions .....	62
Closely comparable trends.....	62
Reporting issues .....	62

Improving the annual crime measure: weighting and significance testing .....	63
<b>Appendix A Chapter 2: additional tables .....</b>	<b>64</b>
<b>Appendix B Chapter 2: the impact of increasing the sample size on PFA level estimates</b>	<b>67</b>
<b>Appendix C Chapter 2: additional information on analysis of the impact of changes to the degree of clustering.....</b>	<b>70</b>
<b>Appendix D Chapter 3: additional tables .....</b>	<b>72</b>
<b>Appendix E Chapter 4: additional tables .....</b>	<b>73</b>
<b>Appendix F Chapter 4: issues for an alternative measure of prevalence rates .....</b>	<b>75</b>

## Tables

Table 2.1	Main features of the BCS core sample design since 1996.....	15
Table 2.2	Design factors attributable to disproportionate selection probabilities for national estimates.....	17
Table 2.3	Design factors attributable to disproportionate selection probabilities, by inner city.....	18
Table 2.4	Average rho and design factors by level of clustering.....	20
Table 2.5	Average rho and design factors for inner city and other areas, by level of clustering .....	21
Table 2.6	Average rho and design factors for inner city and other areas, by level of clustering .....	21
Table 3.1	Comparison of BCS response rates with other social surveys (2007-08) (England only) .....	38
Table 4.1	Summary of how incidents were assigned to months (all incidents 2001/2 to 2008/9) .....	49

## Figures

Figure 3.1	Prevalence of all crime, personal crime and household crime .....	25
Figure 3.2	All household crime, all personal crime and all violence - prevalence by age (of respondent) .....	25
Figure 3.3	All violence - prevalence by age within sex .....	27
Figure 3.4	All violence - prevalence by age within inner city .....	28
Figure 3.5	All violence - prevalence by age within 10% most deprived areas (England only) .....	29
Figure 3.6	All violence - prevalence by age within 'Hard pressed' ACORN category .....	29
Figure 3.7	All violence - incidence by age within 10% most deprived areas (England only) .....	30
Figure 3.8	All violence - incidence by age within 'Hard Pressed' ACORN category .....	30
Figure 3.9	Estimated response rates by age within sex .....	33
Figure 3.10	Estimated response rates by individual age vs. incidence of violence (males) .....	33
Figure 3.11	Estimated response rates by individual age vs. incidence of violence (females) .....	34
Figure 3.12	Estimated response rates by broad age for 10% most deprived areas versus other areas (males) .....	35
Figure 3.13	Estimated response rates by broad age for 10% most deprived areas versus other areas (females) .....	36
Figure 4.1	Example of interview months and crime reference periods for interviews between April 2007 and March 2009.....	46
Figure 4.2	Comparison of annual and interview measures using simulated data – rise and fall in crime levels .....	52
Figure 4.3	Comparison of annual and interview measures using simulated data – Y1 to Y4 .....	53
Figure 4.4	Comparison of annual and interview measures using simulated data – Y1+2Q to Y4+2Q.....	53
Figure 4.5	Vandalism – trends for annual crime and years of interviews (rates per 10,000).....	56
Figure 4.6	Burglary – trends for annual crime and years of interviews (rates per 10,000) .....	56
Figure 4.7	Burglary with entry – trends for annual crime and years of interviews (rates per 10,000).....	57
Figure 4.8	All vehicle-related theft – trends for annual crime and years of interviews (rates per 10,000) .....	57
Figure 4.9	Theft from vehicles – trends for annual crime and years of interviews (rates per 10,000) .....	58
Figure 4.10	Theft of vehicles – trends for annual crime and years of interviews (rates per 10,000) .....	58
Figure 4.11	All household crime – trends for annual crime and years of interviews (rates per 10,000) .....	59
Figure 4.12	All personal crime – trends for annual crime and years of interviews (rates per 10,000) .....	59
Figure 4.13	Theft from the person – trends for annual crime and years of interviews (rates per 10,000) .....	60
Figure 4.14	All violence – trends for annual crime and years of interviews (rates per 10,000).....	60
Figure 4.15	Violence with injury – trends for annual crime and years of interviews (rates per 10,000).....	61
Figure 4.16	Violence without injury – trends for annual crime and years of interviews (rates per 10,000).....	61
Figure 4.17	Domestic violence – trends for annual crime and years of interviews (rates per 10,000) .....	62

## Acknowledgements

The project was initially managed by Stephen Roe and subsequently by John Flatley, both in the Crime Surveys Section of the Home Office. NatCen's research team acknowledges their help in defining the scope of the methodological project, in the provision of data files and documentation and in comments on drafts of the reports.

Professor Chris Skinner of the School of Social Sciences, University of Southampton, acted as a consultant to the NatCen team, taking part in meetings with the Home Office researchers and advising on the approach to the analysis and on draft reports.

We are very grateful to the BCS Team at TNS-BMRB Social Research who provided certain data files that have not been deposited at the ESDS UK Data Archive. These were required, for example, to enable steps in the weighting to be replicated.

Finally, we acknowledge the participation in the BCS of the tens of thousands of members of the public whose addresses are selected each year.

## Disclaimer

The authors of the report wish to point out that they are responsible for the findings reported here. In particular, the views expressed in this report are not necessarily those of the Home Office or of staff at the Home Office.

# Executive summary

The British Crime Survey (BCS) is one of the largest social research surveys in the UK. It is also one of the surveys that is longest-established, giving a record of trends in victimisation among the adult population on a consistent basis of measurement for almost thirty years.

During this long period, the technical design of the BCS has changed in a number of respects. When these changes have occurred, there have been methodological studies to assess their impact on the survey's findings and their interpretation. On some occasions in the past, methodological studies have preceded the adoption of a significant change in the design, which is the preferred approach if it is feasible.

In this report, three aspects of the recent design of the BCS are examined:

- Changes in the sample design and clustering of addresses, in particular to assess the adoption of an unclustered design in part of the sample from April 2008 onwards,
- Variations in the response rate, specifically looking at whether there is evidence of lower response rates among groups that may be more heavily victimised,
- The adoption of continuous interviewing in 2001 had the consequence that the victimisation experiences reported cover a wider span of dates than in the previous survey design.

The work programmes for these three inquiries had a number of similarities. In each case, a data set was constructed that combined data from a number of BCS sweeps. As a check on the data and the weighting, certain headline BCS analyses were replicated. The analyses all relate to the core sample of adults aged 16 and over resident in households and national estimates of victimisation rather than those for Police Force Areas.

Although the focus of this report is on the BCS, the findings of the methodological review may be of interest to those dealing with other social surveys, as the findings have a bearing on survey methods in general.

## Efficiency of different sample designs

The sample design of the BCS has changed on three occasions since 1998:

- Up to 1998, 'inner city' areas were oversampled, in order to boost the number of victims included in the sample;
- From 2000 onwards, there was a requirement to produce estimates for individual Police Force Areas (PFAs); this essentially produced the reverse of the old design, since more of the sample was in 'low crime' areas, although this was mitigated by an increase in the overall sample size; in 2004 the requirement was increased to 1,000 interviews per PFA, which increased further the share of the sample devoted to more rural parts of England and Wales; however, the overall sample size was increase as well;
- From April 2008, while adhering to the sample design based on PFAs, the sample in areas of high population density was no longer clustered within postcode sectors; in areas of low population density clustering was increased.

The estimation of the impact of these changes was based on the calculation of standard errors, and these have been summarised using design factors and effective sample sizes. These measures indicate how well the sample design performs, relative to a simple random sample<sup>1</sup>.

The analysis looked in turn at the implications of varying selection probabilities, the degree of clustering, the rate of victimisation and the overall sample size. The effective sample size has been greatest, as expected, when the sample has been more unclustered in areas where victimisation tends to occur more often.

The requirement to provide estimates for PFAs reduced sample efficiency, particularly from 2004 when the minimum sample was 1,000 interviews per PFA; however, we show that disproportionate selection probabilities have had only a modest impact on the precision of crime estimates.

In common with most social surveys, in most years the BCS used a clustered multi-stage sample design. Clustering tends to increase the homogeneity of respondents within a cluster. This is measured by the intra-class correlation coefficient ( $\rho$ ). As theory suggests, the partially clustered design adopted in April 2008 had a higher efficiency overall (95%) than previous designs based on whole sectors (91%) or quarter-sectors (90%).

These analyses show that under all designs the BCS has generated estimates of victimisation with low levels of variance. Changes in the sample design have not affected the ability of the BCS to identify trends in victimisation. However, the most recent partly-clustered design is shown to provide the most efficient sample.

## Response rates among young people in 'high crime' areas

This part of the review was concerned with the possibility that BCS estimates are biased. It has been suggested that lower response rates among young people in inner city areas have led to under-estimates of victimisation. It is not suggested that such bias affects the validity of trends over time, since the bias suggested could well be constant. It is worth noting here that the BCS has maintained a relatively high response rate of around 75 per cent over the last decade whilst many other surveys have experienced a decline.

In this part of the work, data-sets from five consecutive years were combined to give a large enough base for analysis of victimisation and response by individual years of age (within sex) and for comparisons by area type. Response rates were also produced for individual years of the survey from 2002/03 onwards.

The analysis confirms that rates of personal victimisation<sup>2</sup> are highest among young people (particularly those aged 16-19), whilst corresponding response rates, particularly amongst young men, are relatively low. Response rates by age were shown to be as low as 55 per cent (for young men aged 19 to 21). Other surveys often exhibit similar patterns of response, but not usually combined with the same degree of correlation with the topic of the study as with the measurement of victimisation.

---

<sup>1</sup> In this discussion, we are not concerned with the fieldwork efficiency of the designs for interviewer workloads, which have implications for survey costs.

<sup>2</sup> The BCS divides victimisation incidents into two broad classes, namely 'personal' incidents and 'household' or 'property-related' incidents; the focus here is solely on personal victimisation.

Calibration weights are used on the BCS to ensure that the achieved sample from which the survey estimates are derived, reflects the profile of the population. As with all weighting adjustments, the underlying assumption is that (within weighting cells) non-respondents are not systematically different from respondents in respect of the survey measures i.e. they experience similar victimisation rates to those who are interviewed.

A comparison of the profile of the eligible sample with that of weighted achieved sample showed that calibration weighting brings the achieved sample largely into line with the “correct” profile. However, the adjustments do leave some types of area and types of dwelling slightly under-represented. This suggests that it might be possible to increase the precision of estimates by including direct non-response adjustments in the weighting scheme. For example, details of the accommodation are observed by the interviewer, regardless of whether an interview is achieved, so adjustments could be made for differential non-response by accommodation type. However, the impact of such refinements to the weighting would be quite minor and almost certainly within sampling error; therefore it is not recommended that estimates for past years are re-published.

A major constraint on the analysis was the lack of information on the characteristics of non-respondents. For example, response rates for young people could only be estimated indirectly by reference to external population estimates. Another indirect method used was to estimate the extent to which non-respondents would *need* to be victimised for the omission of these incidents to have a statistically significant effect on the overall BCS estimates. This showed that non-respondents would have to experience at least 25 per cent more (or less) victimisation than respondents for BCS estimates to fall outside of the confidence interval around current estimates.

Although the small amount of available evidence suggests that non-respondents to the BCS are *not* more highly victimised, it is nevertheless conceivable that they could experience such a relative degree of (violent) victimisation compared to respondents. A statistical model of the likelihood of violent victimisation suggested that lifestyle characteristics such as frequenting pub/clubs increase the overall risk of experiencing violent crime and, if such characteristics are also related to the probability of response, this raises the possibility that non-respondents (particularly in younger age groups) might be more highly victimised than their responding counterparts.

Despite this possibility, our conclusion is that there is no compelling evidence to suggest that the survey results are significantly biased due to the under-representation of young people in certain types of area.

## Reporting period and the data collection period

Prior to 2001, the reports published on the BCS used data collected in a specific fieldwork period and experienced during the previous calendar year. For example, the report on the BCS in 2000 used interviews carried out between January and May 2000 and reported on experiences of victimisation in 1999. With the move to continuous interviewing from April 2001, annual reports continued to use interviews completed in a 12 month period (and quarterly updates use a rolling 12-month fieldwork period). However, the victimisation experiences span 23 months, centred on March.

This approach allows headline findings to be published promptly. The main annual report released in July deals with interviews carried out in the fieldwork year April to March. Were it decided to wait until all victimisation for this period had been recorded, researchers would have to wait an additional year for complete datasets to be available for analysis.

Using a dataset consisting of BCS core sample interviews from April 2001 to March 2009, each incident was assigned to a specific month<sup>3</sup>. This allowed monthly rates to be summed for a number of distinct periods. It was then possible to plot the estimates for the two series for a number of aggregate types of crime, such as 'all vandalism', 'all household crime' and 'all violence'. In every case, the two measures follow an almost identical trajectory. Where there were slight deviations, they could be accounted for by examining the month-by month estimates. In short, there was no evidence that the different basis for reporting would have produced different findings in the period 2001-2009. However, it has to be noted that there was a steady decline in the incidence of most crime types over this period.

The review also considered the potential for producing estimates of the prevalence of victimisation, that is the number of individuals experiencing one or more incidents in a 12 month period. Repeat victimisation presents difficulties in the production of such estimates. The report considers the potential of time series modelling or matching approaches to enable such estimates to be produced.

---

<sup>3</sup> Respondents stated the month for 64 per cent; and the quarter for 36 per cent of incidents ('victim forms'. The latter cases were assigned at random to a specific month. Only 711 of the 143,976 incidents had to be assigned to a month within a year.



# 1 Introduction

The British Crime Survey is one of the major social research projects conducted in the UK. In 2009-10, the specification required a representative sample of 47,000 adults (aged 16 and over) resident in private households, plus a boost sample of 4,000 respondents aged 16 to 24. From January 2009, the sample has also included a representative sample of 4,000 young men and women aged 10 to 15.

The BCS is used to provide estimates of victimisation and to measure a number of Home Office performance targets, both nationally and for individual Police Force Areas<sup>4</sup>.

The series began in 1982 and certain features of the questionnaire have remained virtually unchanged. As a result, the BCS has a remarkable record of providing an unbroken trend line over the period of nearly 30 years in which the study has been carried out. During the BCS, the trend in victimisation was steadily upward until 1995 (the BCS interviews conducted in 1996), when victimisation prevalence peaked at nearly 40 per cent of respondents. Since then, the trend has been downward, initially quite steeply but more recently changing little from year to year. Over the latter half of the period, the rate of victimisation has fallen to around 20 per cent of adults in the sample

## 1.1 Previous methodological studies related to the BCS

Over the period during which BCS has been conducted a number of technical changes have been adopted. For example, prior to 1992 the Electoral Register had been used as the sample frame. In 1992 the Postcode Address File (PAF) was adopted and the methodological implications of that change were assessed (Lynn, 1992: Assessment of the impact of changing the sampling frame for the British Crime Survey, Home Office, 1992). Another major change was the adoption of continuous interviewing in 2001. The change was informed by the Review of Methodology conducted by Peter Lynn and David Elliott, published in 2000. (Lynn, P. and D. Elliot, *The British Crime Survey: A Review of Methodology*, National Centre for Social Research, London: Home Office, 2000.)

Much of the methodological research cited above was concerned with sample design. In addition, a qualitative study was carried out following-up respondents to the 1996 BCS, which investigated some aspects of the reporting of victimisation incidents. The main issues were around recall – whether respondents could remember incidents - and sensitivity – whether, having recalled an incident, the respondent would be willing to discuss it with an interviewer. This study involved in-depth interviews with 35 individuals who had been interviewed on the BCS some weeks prior to being approached to take part in the qualitative study. This was reported by White, C. and Lewis J. (1998) Following up the British Crime Survey 1996 - a qualitative study. SCPR report. London: Social and Community Planning Research.

Another aspect of the ability of respondents to recall incidents is the length of the reference period in questions. Up to 2000, the BCS interviews were conducted in the first few months of the year, with respondents being asked to report any incident they had experienced since the start of January in the previous year. With the change to continuous data collection in 2001, the reference period was to the start of the month in which the interview was taking place, and so varied from 12 to 13 months. For

---

<sup>4</sup> Since 2002 the BCS has been the main source used by the Home Office to measure progress against two of its Public Service Agreement targets. In addition, the BCS is used as part of the Police Performance and Assessment Framework to measure individual forces progress against a number of Statutory Performance

part of the BCS survey in 2001, both the old and new survey designs were used in parallel, so that victimisation estimates based on both approaches could be compared. In the event, relatively few incidents occur at the margins of the reference period, so the two approaches did not produce substantially different estimates. The results were reported by Peter Lynn in a presentation to a conference at the Royal Statistical Society in 2002.

Peter Lynn also reported on an experiment carried out on part of the 1996 BCS. Where a sample member had been selected but had then refused to take part in the full interview, the interviewer attempted to establish whether the respondent had experienced one of a short list of incidents. The exercise showed no evidence that non-respondents were any more likely to have experienced these types of incidents than those who were interviewed (Lynn, P (1997) *Collecting data on non-respondents to the British Crime Survey*. Unpublished report to the Home Office. London: Social and Community Planning Research.)<sup>5</sup>

The BCS was one of the twelve major surveys included in the Census-linked Study of Survey Non-Response (Freeth, 2005). The July and August 2001 samples for the BCS were matched to data from the 2001 Census. Non-imputed Census records were linked successfully to 81.6% of responding (not all of which gave permission for the linkage) and 86.9% of non-responding BCS sample households. One significant limitation of such studies is that survey interviewers usually do not record age and sex details of individuals who refused. This means that the analysis had to be concerned with household characteristics and some details of the individual identified as the Household Reference Person.

The most recent methodological study conducted on the BCS was a review of the options for extending the coverage of the sample to include children and people living in communal establishments (Kevin Pickering, Patten Smith, Caroline Bryson and Christine Farmer (2008) Home Office Research Report 6).

Of course, the BCS is one of the many studies of victimisation that are carried out in many countries using broadly similar methods. The original design of the BCS owed a considerable amount to methodological work carried out in the United States for the National Victimization Survey<sup>6</sup>. A number of European and other countries now carry out regular surveys of victimisation, and many of the organisations responsible for these surveys have conducted methodological research. There is an International Survey of Victimization and surveys of Commercial victimisation. In the UK, a specialised study of 'elder abuse'<sup>7</sup>. was carried out to estimate the extent to which there may be victimisation of vulnerable elderly people that may be concealed from standard survey data collection approaches; this adopted essentially similar methods of research to those developed for the estimation of victimisation in the wider population. In short, there is a considerable body of experience on which

---

indicators. This was the driving force behind the minimum sample size per Police Force Area. (See technical report for more details.)

<sup>5</sup> In 2005, Steven Hope reported on a methodological study of the Scottish Crime Survey, which also showed that non-respondents were *less* likely to take part in the survey than those who had been victims

<sup>6</sup> A major review of criminal justice statistics, including the allocation of responsibility for their collection and dissemination, was conducted in the United States in 2007-09, concerned with statistics collected by official agencies in the Criminal Justice System as well as covering arrangements for collecting information on victimisation by household survey methods (Groves, *et al*, 2009, available for download or viewing at: [http://www.nap.edu/openbook.php?record\\_id=12671&page=R1](http://www.nap.edu/openbook.php?record_id=12671&page=R1)).

<sup>7</sup> O'Keeffe, M, Hills, A, Doyle, M. McCreddie, C., Scholes, S., Constantine, R, Tinker, A, Manthorpe, J., Biggs, S., Erens, R. (2007) UK Study of Abuse and Neglect of Older People :Prevalence Survey Report, National Centre for Social Research

those responsible for management of victimisation surveys may draw for information on technical aspects of survey design.

## 1.2 The BCS Methods Review, 2009-10

The methodological research project documented in this report examined three aspects of the British Crime Survey, England and Wales. These are outlined briefly in this section.

### Sampling

Previous sweeps of the BCS were carried out in 1982, 1984, 1988, 1992, 1994, 1996, 1998, 2000 and it has been as a continuous survey from 2001 onwards. From 2000 onwards, the original practice of boosting the sample in 'inner city' areas was replaced with a sample design based on police force areas (PFAs). For most of the period since 2000, the target has been 1,000 interviews per PFA, with larger numbers in a few very large PFAs. While this design has been retained, the degree of clustering was changed from April 2008. In areas of high population density, the sample was no longer clustered. This was compensated by an increase in clustering in areas of low population density.

The aim of the first part of the project is to review the impact of the changes in sample design since 1998 on the reliability of BCS crime trend estimates.

### Response rate

The BCS response rate has been maintained at a level of around 75 per cent during the last decade when response rates on many other major government-sponsored studies have fallen. Calibration weighting ensures that the achieved sample is consistent with the structure of the population by sex and age within Government Office Regions. However, concern has been expressed that the response rate is lower among certain age categories (such as young people aged 16 to 24) and in certain areas, such as areas of social deprivation where crime rates may be relatively high. If the experience of the people not interviewed was very different from that of respondents, that could result in a degree of bias.

### Reporting period and reference period

The data reported in annual and quarterly reports are based on interviewing in a 12 month period. However, in each of the 12 months of fieldwork, the reference period is up to 13 months. As a result of this, the victimisation experiences reported by respondents interviewed in a period of 12 months relate to a broader span of 23 months. The aim is to examine the implications for crime trends of reporting on data *collected* during a financial year compared to reporting on prevalence or incidence of crime *experienced* during a financial year.

## 1.3 Approach to conducting the methodological research

For all three of these strands of methodological work, the approach has been based on developing a data file covering a series of BCS years. The reasons for doing this varied according to the strand of the methodological research. In the case of the response rate analysis it was to ensure the base for analysis was large enough to allow disaggregation by age, sex and type of area. For the other two strands, the combined dataset allowed direct comparisons to be made between the data collected at different periods under varying research designs.

As well as covering productive interviews, it was necessary to examine records of the original samples of addresses. It was also necessary to gain a detailed appreciation of the way in which the BCS datasets have been weighted in preparation for estimates to be compiled.

The main published BCS estimates of crime are based on analysis of the core sample, which is a nationally representative sample of persons aged 16 and over living in private households. Over the years the BCS has included a number of additional boost samples, including ethnic minorities in some years. However, the boost samples have been excluded from the analyses covered in this report, in line with BCS procedures for producing routine annual estimates.

## 2 Evaluating changes in the BCS sample design

This section assesses the impact of changes in sample design since 1996 on the reliability of BCS crime trend estimates. The year 1996 is chosen as the starting point because it was the year in which victimisation rates peaked and it represents the penultimate year in which the 'inner city' sample was boosted.

### 2.1 Changes to the BCS sample design

The BCS sample design has undergone a number of changes since 1996, which include:

- changes in the selection probabilities for sampling of areas,
- changes in the degree to which the sample is clustered,
- and an increase in sample size.

These changes are outlined in more detail below.

The analysis in this section is focused on the core BCS sample, which is a nationally representative sample of persons aged 16 and over resident in households. Over the years the BCS has included a number of additional boost samples. However, these have been excluded from this analysis as the main published BCS estimates of crime are based solely on the core sample.

#### Disproportionate sampling

From the first BCS in 1982 until 1998 the BCS selected addresses in inner city areas at twice the rate of the rest of the core sample. The intention was to over-sample areas known to have generally higher incidence of crime. The sample size during these years was smaller than more recently, and in this context the over-sampling of 'high-crime' areas a sensible option, since victimisation was a relatively rare experience. Whilst this feature of the design was dropped in 2000, the increases in overall sample size means the number of interviews carried out in inner city areas has remained fairly constant, despite them making up a smaller proportion of the overall sample. Table A.1 in Appendix A shows the core sample size by inner city and other areas. The sample size in inner city areas was 3,870 in 1996 (24% of the overall sample) and 3,629 in 2008-9 (8% of the overall sample size).

Since 2000 (and at the time of drafting in 2010), there has been a requirement to attain a minimum sample size in each of the 42 Police Force Areas (PFAs) and as a result less populous PFAs have been over-sampled relative to larger ones. This design was adopted to allow PFA-level estimates to be produced. In 2000 a minimum of 300 interviews were sought in each PFA. When the survey became continuous in 2001/02 the minimum number of interviews per PFA was increased to around 700. In 2004/05 the sample size was increased again to approximately 46,000 interviews per year with the aim of achieving at least 1,000 interviews in each PFA. This remains the current target sample size for the core sample of persons aged 16 and over in 2009/10.

Over-sampling less populous PFAs means areas with higher crime rates are no longer over-sampled. Instead the PFAs that are over-sampled tend to have largely rural and suburban populations, areas which are generally associated with lower crime rates. The more populous PFAs tend to be located in metropolitan areas, which are under-sampled in this design. One of the aims of the analysis was to look at the impact of this change on the precision and level of victimisation incidence rates in inner city areas.

## Clustering

The BCS is a large-scale national survey that uses face to face interviewing. As such the sample is geographically clustered in order to reduce fieldwork costs and assist field-staff management. A number of changes have been made to the level of clustering used in the BCS. From 1996 to 2000 the BCS used quarter-postcode sectors as clusters. From 2001-2 to 2007-8 whole postcode sectors were used. In 2008/09 the sample design was changed in several respects:

- in areas with a high population density no clustering was used,
- in medium population density areas the degree of clustering was similar to previous surveys and
- in low population density areas the degree of clustering was increased as compared to previous waves.

This design marked a return to the notion of making the sample design most efficient in areas of the highest incidence of crime. It was based on the proposition that the additional costs of having an unclustered sample in areas of high population density would be offset by greater clustering in areas of low population density. In effect, there was a trade-off between greater levels of precision for estimates in areas with greater levels of crime and potentially lower precision in areas where levels of crime were lower.

## Sample size

The BCS core sample size has more than trebled since 1996; from 15,000 to the current size of 46,000. If everything else in the survey design is held constant, an increase in sample size will always improve the precision of the survey estimates. A minimum sample size has been set in PFAs, which means the precision of PFA-level estimates will increase, allowing better cross-PFA comparisons to be made (this is discussed further in Appendix B).

The main changes in the BCS sample design are summarised in Table 2.1.

**Table 2.1 Main features of the BCS core sample design since 1996**

Year	Core target sample size	Achieved sample size	Main design features	Clusters
08/09	46,000	46,289	Disproportionate sampling by PFAs to get a min of 1,000 per PFA	Unclustered in areas of high population density, highly clustered in rural areas
07/08	46,000	46,983	Disproportionate sampling by PFAs to get a min of 1,000 per PFA	Whole postcode sectors (32 issued addresses per PSU)
06/07	46,000	47,023	Disproportionate sampling by PFAs to get a min of 1,000 per PFA	Whole postcode sectors (32 issued per PSU, 16 in high density areas)
05/06	46,000	47,796	Disproportionate sampling by PFAs to get a min of 1,000 per PFA	Whole postcode sectors (32 issued per PSU, 16 in high density areas)
04/05	46,000	45,120	Disproportionate sampling by PFAs to get a min of 1,000 per PFA	Whole postcode sectors (32 issued per PSU, 16 in high density areas)
03/04	37,000	37,931	Disproportionate sampling by PFAs to get a min of 600-700 per PFA	Whole postcode sectors (32 issued per PSU, 16 in high density areas)
02/03	37,000	39,249	Disproportionate sampling by PFAs to get a min of 600-700 per PFA	Whole postcode sectors (32 issued per PSU, 16 in high density areas)
01/02	37,000	32,824	Moved to a continuous fieldwork period. Disproportionate sampling by PFAs to get a min of 600-700 per PFA	Whole postcode sectors (32 issued per PSU, 16 in high density areas)
2000	20,000	19,411	Disproportionate sampling by PFAs to get a min of 300 per PFA	Quarter postcode sectors (32 issued per PSU)
1998	15,000	14,947	Inner city areas sampled at twice the rate of other areas	Quarter postcode sectors (36 issued in inner city areas, 32 in other areas)
1996	15,000	16,348	Inner city areas sampled at twice the rate of other areas	Quarter postcode sectors (30 issued in inner city areas, 27 in other areas)

Note: The survey only covers the population aged 16 years and over who are resident in households; persons in institutions such as army barracks, student halls of residence, residential care or prison are excluded. However, the impact of this has been shown to be small<sup>8</sup>.

<sup>8</sup> See Pickering, K., Smith, P., Bryson, C. and Farmer, C. (2008) *British Crime Survey: options for*

A list of the crime measures used in the analysis is given in Table A.2 in Appendix A.

## 2.2 Estimating the effects of changes in sample design

It seems worth introducing the statistical measures used in the analysis before presenting the analysis itself. These measures are all concerned with the precision of survey estimates, making use of a 'simple random sample' as the benchmark design with ideal properties. In practice, unit costs of survey administration are reduced by adopting more complex sample designs that involve clustering and disproportionate sampling, so most surveys use designs that fall some way short of the ideal.

The impact of the sample design on key estimates can be quantified using *design effects* and *design factors*. The *design effect* shows the impact of the sample design on the variance of the survey estimates. It can be used to show how much information has been lost or gained by using a complex sample design, as compared to a simple random sample. The *design factor* is the square root of the design effect. It shows how much the standard errors of the survey estimates have been inflated by the adoption of the actual design by reference to a simple random sample of the same size. The design factors for key estimates under different sample designs are used to show how the precision of crime estimates has been affected by changes in the sample design.

In turn the design factor can be used to generate the *effective sample size*, which is a related measure of sample precision; it shows the size of a simple random sample that would have achieved the same level of precision as the design under consideration. If the effective sample size is close to the actual sample size then the sample design is efficient with a good level of precision. The lower the effective sample size is, the lower the level of precision. The *sample efficiency* is the ratio of the effective sample size to the actual sample size and indicates how close the effective sample size is to the actual sample size.

To illustrate the key concepts, consider the following hypothetical example. A clustered design with a sample of 4,000 cases and a *design factor* of 1.15 is equivalent to a simple random sample of 3,025 cases ( $4000 / (1.15 \times 1.15)$ ).

The impact of the changes described in Section 2.1 are expected to be on the amount of variation in the crime estimates (known as variance), rather than the levels of bias. It is unlikely that such changes in the sample design would affect an individual's propensity to respond (as discussed in Chapter 3), hence we would not expect to find impacts on the levels of bias in estimates. A high level of variance will result in less precise estimates. We could be less certain that the sample estimates were a true reflection of the population. It would also reduce the ability to identify statistically significant changes over time. Any increase, or decrease, in the variance of the crime estimates will be reflected in the size of the standard errors and can be described using design factors and effective sample size.

### The impact of disproportionate selection probabilities

The BCS sample design in 2009-10 selects addresses with unequal selection probabilities across PFAs; that is, addresses in less populous PFAs have a higher chance of being selected. Over-sampling a specific sub-group within a population has obvious advantages for analysis of that particular sub-group. However, unless the over-sampling is associated with a key stratification factor



(e.g. greater probability of victimisation), it will be detrimental to the effective sample size for estimates relating to the whole sample.

With disproportionate sampling, each case receives a weight that is the inverse of the selection probability for the case. Whilst the selection weights correct bias, the more they vary the more they reduce the effective sample size. Less efficient samples have larger standard errors, resulting in wider confidence intervals<sup>9</sup> around the survey estimates; this means there is less certainty that the survey estimate is close to the true population value.

A design in which the population in smaller PFAs is over-sampled relative to that in larger PFAs involves a trade off between a reduction in the precision of national estimates and higher precision for PFA-level estimates. The importance of PFA-level estimates versus national estimates, and the extent to which both are used, need to be taken into account when considering the effects of disproportionate sampling within PFAs. However, over the same period the increase in the overall BCS sample size should have lessened the impact of over-sampling smaller PFAs (and under-sampling large ones) on the precision of national estimates.

Table 2.2 shows the increase in the design factors that is attributable to the disproportionate selection probabilities. This has been estimated directly from the selection weights<sup>10</sup>. Other aspects of the sample design, such as clustering, have been ignored in order to focus on the impact of the disproportionate selection probabilities, although the same patterns occur when this information is included (see Table A.3 in Appendix A).

Table 2.2 Design factors attributable to disproportionate selection probabilities for national estimates				
BCS year	Actual sample size	Effective sample size	Sample efficiency	Design factors of the core sample
2008-09	44,002 <sup>†</sup>	38,167	87%	1.07
2007-08	46,983	39,484	84%	1.09
2006-07	47,203	40,221	85%	1.08
2005-06	47,796	40,833	85%	1.08
2004-05	45,120	38,967	86%	1.08
2003-04	37,931	34,587	91%	1.05
2002-03	36,479	33,811	93%	1.04
2000	19,410	18,992	98%	1.01
1998	14,947	13,958	93%	1.04
1996	16,348	15,363	94%	1.03

<sup>†</sup> Each survey year contains a small number of cases rolled forward from the previous one. There were 2,284 cases in the 2008-9 survey data that had been rolled forward from 2007-08. These were excluded for the purposes of evaluating the effects of the sample design as they would have been selected under the previous design.

<sup>9</sup> A 95% confidence interval is constructed in such a way that 95 times out of 100 it captures the true population value of the survey estimate. A narrow interval suggests a better level of precision as it suggests the survey estimate is closer to the actual population value.

<sup>10</sup> Estimated as the actual sample size divided by the effective sample size, where the effective sample size is estimated as the sum of the weights squared divided by the sum of the squared weights;  $neff = (\sum n_i w_i)^2 / \sum (n_i w_i^2)$ .

Sample efficiency decreases (and design factors increase) when selection probabilities become more unequal. In 1996 and 1998, addresses in inner city areas had selection probabilities twice as high as other areas. The resulting design factors are 1.03 and 1.04, respectively, as shown in Table 2.2. In 2000 the inner city sampling was removed and instead less populous PFAs were over-sampled relative to larger ones to give a minimum sample size of 300 interviews per PFA. This was a more efficient design with less variable address selection weights; in effect, each PFA received its proportionate share of the national population, with only a small increase in selection for the smallest PFAs. In 2001 the numbers per PFA increased. Although the overall sample size also increased, differences in the size of PFAs meant address selection probabilities needed to be more unequal in order to ensure the minimum requirement could be met and the design factors increased accordingly. The design factors increase again in 2004-5 when the sample design moved to a minimum of 1,000 interviews per PFA.

Sample efficiency is higher for sub-group estimates for groups identified by variables used to stratify the sample (such as PFA-level estimates for BCS years when PFA was used as the main stratification variable). This is because the selection probabilities will be more equal within these groups. Likewise, sample efficiency will decrease for sub-groups that overlap with variables used to stratify the sample. The sample efficiency of estimates in inner city areas is close to one in 1996. It drops from 2000 onwards, when this variable was no longer used to stratify the sample. The sample efficiency appears to increase slightly in 2008-9. It may be that the changes in sample design have made the selection probabilities less unequal, or it could simply be due to random fluctuations. More years of data would be required to say for certain. This is shown in Table 2.3.

Table 2.3 Design factors attributable to disproportionate selection probabilities, by inner city								
BCS year	Inner city areas				Other core sample			
	Actual sample size	Effective sample size	Sample efficiency	Design factors of the core sample	Actual sample size	Effective sample size	Sample efficiency	Design factors of the core sample
2008-09	3,629	3,359	93%	1.04	40,373	34,851	86%	1.08
2007-08	3,699	3,237	88%	1.07	43,284	36,347	84%	1.09
2006-07	3,675	3,323	90%	1.05	43,528	36,986	85%	1.08
2005-06	3,736	3,432	92%	1.04	44,060	37,452	85%	1.08
2004-05	3,641	3,329	91%	1.05	41,479	35,705	86%	1.08
2003-04	3,395	3,219	95%	1.03	34,536	31,401	91%	1.05
2002-03	3,219	3,070	95%	1.02	33,260	30,754	92%	1.04
2000	2,363	2,322	98%	1.01	17,047	16,670	98%	1.01
1998	3,542	3,542	100%	1.00	11,405	11,405	100%	1.00
1996	3,870	3,869	100%	1.00	12,478	12,478	100%	1.00

The changes in sample design between 1996 and 2000 made PFA level estimates more efficient, since this variable is used as a stratification variable and the selection probabilities will be equal for the address within PFAs. The design factors for PFAs in 1996 ranged between 1.00 and 1.07. In all other years they are between 1.00 and 1.01 (not shown).

The PFA-level crime estimates benefit from the changes in sample design and the resulting increase in sample size per PFA. The sample design does not affect PFA-level estimates because the selection probabilities are equal within each PFA; the design factors within PFAs are close to one. However, the sample design will have an impact on national crime estimates, which is shown in Table A.4 in the Appendix. This table gives the standard errors and confidence intervals for separate crime estimates, with and without stratification. The difference between the two confidence intervals is generally small. For example, vandalism in 2008-9 had a prevalence of 1,161 per 10,000. Confidence intervals that took disproportionate sampling into account range from 1,093 to 1,277; confidence intervals that assume simple random sampling are 1,099 to 1,223. The disproportionate selection probabilities do not have a large impact on the precision of national crime estimates.

## The impact of changes in the degree of clustering

This section shows the impact of variations in cluster size on the precision of crime estimates, and therefore the precision of crime trends. Of particular interest is the impact on inner city areas and how the recent changes to the sample design have affected these areas. It must be noted that there are many important and competing factors that influence sample design, such as survey costs and interviewer management. This section looks at the effects of clustering in isolation and excludes other factors that would have been considered at the design stage.

Prior to 2008-9 the BCS used a clustered multi-stage sample design, similar in many respects to most other major national surveys in England and Wales. At the first stage, PSUs consisting of geographic areas (these were individual postcode sectors or groups of smaller postcode sectors) were selected. At the second stage, addresses were selected from each PSU. The resulting sample was clustered in geographical areas, with each cluster usually being an assignment for an interviewer. From 1996 to 2009 the sample was around 30 - 32 addresses issued per PSU, although from 2001 the sample has been 16 per PSU in 'high density' PSUs (with double the number of PSUs in 'high density' areas).

The number of addresses within each cluster can impact on the precision of the estimates. For a given sample size, the standard error will generally increase as the sample size in each cluster increases. This is because the individuals within a cluster tend to have similar social characteristics and they also tend to be different from individuals in other clusters. This means that clustering may increase the risk of selecting a sample that is different from the population. This feature results in increased width of confidence intervals for survey estimates.

In 2008-9 the BCS moved to a new design with clustering in some areas and no clustering in others, in order to maximize both fieldwork and sample efficiency. While clustering was removed in the most densely populated areas, it was increased in low density rural areas where lower Super Output Areas (SOAs)<sup>11</sup> are used as PSUs<sup>12</sup>. The reduction of clustering in the high population density PSUs is expected to have improved sample efficiency, as this is where a disproportionate amount of victimisation is reported. The increase in clustering in rural areas is expected to have had little impact on victimisation estimates as they tend to contribute much less to the national crime count.

The effects of clustering has been examined using the *intra class correlation coefficient* ( $\rho$ ), which measures the degree of similarity among respondents within a cluster using the crime measures listed

---

<sup>11</sup> A postcode sector contains on average 2,700 addresses. A Lower Super Output Area contains an average of 600 addresses and a Middle Super Output Area contains an average of 3,000.

<sup>12</sup> See the BCS 2008-09 technical report for full details on the sample design:  
<http://www.homeoffice.gov.uk/rds/bcs-methodological.html>

in Table A.2. Rho has values ranging between zero and one; a value of one implies that all the respondents within the cluster are identical, whereas a value of zero means there are no similarities and all the respondents are unique. Respondents are expected to be more alike when clusters are small (in population, not necessarily geographically in the most rural areas). Consequently we would expect highly clustered samples to have larger values of rho. More details about the methods used are given in Appendix C.

Rho was estimated for each of the key crime estimates in Table A.2 for all years of BCS from 1996 onwards. Table 2.4 shows the values of rho by level of clustering averaged over years and variables. The values are generally low (the figures in Table 2.4 have been given to a large number of decimal places because the differences are so small). A design factor of 1.054 means the standard errors for an estimate from a sample using quarter postcode sectors were increased, on average, by 5.4%.

Table 2.4 Average rho and design factors by level of clustering			
PSU	rho	Design factors	Sample efficiency
Quarter sectors (1996, 1998, 2000)	0.0053	1.054	90%
Whole sectors (2001-8)	0.0048	1.046	91%
Partially clustered (2008-9)	0.0037	1.024	95%

As expected, the impact of clustering was higher for samples with quarter sectors (1996, 1998 and 2000) than samples that used whole postcode sectors (2001-8), although the difference between the two was very small. This is reflected by the latter having smaller values of rho and smaller design factors. The partially clustered design has the lowest values overall, despite the increased clustering and smaller PSUs in areas of low population density.

The sample efficiency shows how efficient the clustered sample is in relation to a simple random sample. The partially clustered sample has the most efficient design, losing 5% efficiency due to clustering. The difference between whole and quarter sectors is small, with the two samples losing 10% and 9% efficiency, respectively.

The analysis was repeated separately for inner city and other areas, the results of which are given in Table 2.5. For both inner city and other areas the design factors are highest, and sample efficiency lowest, for the design that used quarter sectors.

The partially clustered design is un-clustered in areas with high population density. Although this does not exactly match the inner-city definition, the overlap is sufficient for the design factors in these areas to be effectively equal to one, giving a sample efficiency of 100%.

The remainder of the partially clustered sample is clustered either within middle or lower SOAs. Here the values of rho are slightly lower than those for the design that used whole sectors, but the two are very close. Of the two designs, the partially clustered sample has a slightly smaller average cluster size (an average of 18 achieved interviews per PSU, compared to 21), which increases the difference in the design factors, although they are still very close.

**Table 2.5** Average rho and design factors for inner city and other areas, by level of clustering

PSU	rho	Design factors	Sample efficiency
<i>Inner city</i>			
Quarter sectors	0.0052	1.050	91%
Whole sectors	0.0046	1.034	94%
Partially clustered	0.0001	1.000	100%
<i>Other areas</i>			
Quarter sectors	0.0053	1.055	90%
Whole sectors	0.0039	1.038	93%
Partially clustered	0.0036	1.030	94%

The analysis was repeated using data for 2008-9 split by main sampling stratum. This shows the difference in precision for the different areas due to the different levels of clustering. The results are given in Table 2.6.

**Table 2.6** Average rho and design factors for inner city and other areas, by level of clustering

PSU	rho	Design factors	Sample efficiency
Un-clustered (cities)	0.000	1.000	100%
Moderately clustered (suburbs)	0.004	1.042	92%
Highly clustered (rural areas)	0.004	1.020	96%

The un-clustered stratum has a rho of zero and sample efficiency of 100%. The values of rho in the moderately and highly clustered strata are very similar; however the highly clustered sample stratum has smaller design effects because the PSUs in these areas had a smaller achieved sample size – 11 interviews on average, compared to 21 in moderately clustered areas. This makes the overall sample efficiency higher in the highly clustered areas.

## 2.3 Summary of findings

In this section we looked at the effects of stratification, clustering and increased sample size on the precision of crime estimates. The overall impact of these changes of sample design on national and PFA level BCS crime trends appears to be small.

Over-sampling smaller PFAs has reduced the precision of overall estimates but this has been mitigated by the increased sample size, as was intended. The selection probabilities became more unequal when the minimum sample size per PFA was increased from 700 to 1000 in 2004-5. However, the further increase in the overall sample size meant the effective sample size for national estimates remained about the same and the precision of crime estimates was not seen to drop. The

resulting design allows more precise estimates to be derived at PFA level. It also allows more detailed analyses to be carried out at PFA level and in general.

The design effects attributable to clustering are generally low, although the movement from quarter sectors to whole sectors as PSUs did bring about a modest improvement in precision. The semi-clustered design for 2008-9 sees further improvements, both for the overall sample but particularly for inner-city areas. The impact of increasing the clustering in rural areas was offset by reducing the sample size per PSU, which reduces the effects of clustering on precision within these areas.

## 3 Response rate for young people in 'high-crime' areas

The aim of the second strand of the review is to examine the hypothesis that there exists a highly victimised group of young people in inner city areas who are under-represented in the BCS to the extent that there is significant bias in overall estimated levels of victimisation.

### 3.1 Background and approach

#### Background

The BCS benefits from a relatively high aggregate response rate compared with other large-scale random probability household surveys in the UK; this has been maintained at around 75 per cent for the past decade. However, in common with other such surveys it suffers from lower levels of response amongst younger people and in certain areas including inner city areas (where access is more difficult for interviewers and co-operation rates are also lower). This is a particularly pertinent issue for the BCS, as personal victimisation is concentrated amongst young people and is generally more prevalent in inner cities. This has led to the accusation that the BCS is biased. This is the issue with which this paper is concerned.

In order to identify the scale of any potential bias we would ideally like to examine response rates by age/sex within different types of area and, correspondingly, estimate the extent to which victimisation is more prevalent amongst young people in inner cities. Were we able to do this reliably we could gauge whether the hypothesis has any real validity. However, there are difficulties with attaining each of these aims.

First, whilst it is straightforward to calculate overall response rates for different area types, it is not possible to directly measure response rates by age group as we cannot identify, using the available data, the age or sex of any selected person who refused or was unable to take part. Instead we must estimate response rates for young people by comparing the profile of the sample with population estimates. That this can only be done accurately at the overall (population) level rather than within smaller areas (such as inner cities) is a serious drawback.

Second, the only reliable data on victimisation we have available is from the BCS itself. Therefore, in order to examine the extent to which crime is more prevalent in certain areas and/or amongst particular groups, we must use data from the BCS. As we have no information on victimisation amongst those who do not respond, we will never be able to prove conclusively whether the BCS is biased or not.

#### Approach

Despite these apparent limitations, we have a wealth of data from past years of the BCS on which to draw. Since 2001/02, the BCS survey design, response rate and overall reported levels of victimisation have remained fairly stable. We have therefore felt able to combine raw data from a number of consecutive survey years to increase the base for young people and for inner city areas up to levels which can support detailed and robust analysis. (More detail is available in Appendix D, Table

D1.) Using these combined datasets we have employed a number of different approaches to examine the hypothesis under consideration.

First of all we have calculated aggregate level crime rates by age and area type in order to understand the extent to which victimisation is more prevalent amongst young people and in certain types of area. We have then looked at response rates by age and area type, with particular focus on estimating response rates for young people aged 16-29. By examining estimated response rates side by side with corresponding victimisation rates we were able to evaluate the potential for bias in the survey estimates.

Next we have examined the extent to which calibration weighting adjusts for different types of non-response, particularly area-level non-response. If we assume that non-respondents of a particular age living in a given area are no more likely to suffer victimisation than respondents with the same characteristics, then the only source of bias would be that which the weighting scheme is unable to adjust for. This assumption is key: if it is false then we can only speculate about the true extent of any bias.

Whilst we cannot prove or disprove this assumption, we can go some way towards testing it by investigating the characteristics and behaviours that are associated with victimisation. If these are also correlates of non-response then this would provide some evidence of the potential for bias. We have done this by modelling the likelihood of being a victim of violent crime. In particular we look at lifestyle characteristics such as the frequency of visits to pub/clubs to see the extent to which these are associated with high levels of victimisation.

Finally we have attempted to estimate the degree to which non-respondents would have to be more highly victimised for there to be a (statistically) significant downward bias in the annual estimates of violent crime. As noted earlier, the lack of information on non-respondents means that we cannot prove conclusively whether or not the BCS is biased: this analysis is intended to put the contention of bias into some kind of context.

We begin by looking at aggregate levels of crime by age to see the degree to which young people are disproportionately the victims of different types of crime.

## 3.2 Broad patterns of victimisation on the BCS

### Overall levels of crime


The BCS divides victimisation incidents into two broad categories: household and personal crimes: adding these together gives an estimate of total crime. In any given year, the BCS estimates that around a quarter of people are affected by crime of some sort with around six or seven per cent being victims of personal crime<sup>13</sup>. This is shown in Figure 3.1 below, which also illustrates that the prevalence of crime has remained fairly stable over the past seven years, perhaps with a small downward trend.

---

<sup>13</sup> Excluding sex offences.




**Figure 3.1** Prevalence of all crime, personal crime and household crime



## Young people and crime

Whilst household crime affects all age groups, personal crime is highly concentrated amongst young people as Figure 3.2 illustrates. It shows prevalence rates for household crime, personal crime and violence (a subset of personal crime) for 5-year age groups pooled over five years of survey data (2003/04 to 2007/08). The dashed lines show the overall prevalence rates for the same period: 19.8% for household crime, 6.5% for personal crime and 3.6% for violence.

**Figure 3.2** All household crime, all personal crime and all violence - prevalence by age (of respondent)



Between the ages of 16 and 49, the likelihood of a BCS respondent being a victim of household crime remains fairly constant<sup>14</sup>. By contrast, personal crime, and violence in particular, clearly has a disproportionate impact on young people. For those aged 16-19, the prevalence of personal crime is three times the overall rate and the prevalence of violence is more than 3.5 times the average.

Given that we wish to study the potential for bias caused by the under-representation of young people, it would seem sensible to concentrate on personal crime rather than household crime. However, all personal crime is probably too general and insensitive for our purposes. For example, a sizeable element (around a quarter) of personal crime incidents are categorised as “other theft of personal property” (as distinct from “theft from person” and “robbery”). Analysis shows that this type of crime, whilst still mainly affecting young people, is just as prevalent in affluent areas as it is in deprived/inner city areas. We could choose to focus instead on, individual crimes such as mugging but these might be too specific to particular types of people/areas. We have therefore chosen to focus on violence (this includes domestic violence, mugging, stranger violence and acquaintance violence) as this is clearly the main driver of personal crime but is not too narrow in its definition.

Looking at prevalence of violence by age, it falls steadily across the younger age groups. However, it is not until ages 35-39 that prevalence falls to the level of the overall average. The implication of this for examining crime affecting “young people” is that we must be careful how we define a young person. The key age group is probably those aged 16-24, with 16-19 year olds being of particular interest. However, it may be wise to extend our definition of young people to include those aged 25-29, if not those aged 30-34 as well.


## Prevalence of violence by age within sex

Figure 3.3 shows the prevalence of violence broken down for males and females separately. Although both men and women aged 16 to 24 are disproportionately affected by violence, it is young men who are most likely to be victims, particularly those aged 16-19. The gap between men and women remains until ages 50-54; thereafter the likelihood of victimisation is almost the same regardless of sex.

---

<sup>14</sup> Individual respondents are asked about household victimisation, which may have affected property owned or used by any household member.

**Figure 3.3** All violence - prevalence by age within sex




### 3.3 Victimisation by age and area type

#### Prevalence of violence by age and area type

We have seen that young people, particularly young men aged 16-24 are those most likely to be the victims of personal crime, particularly crimes involving violence. However, the hypothesis under examination suggests that the problem is particularly acute in “inner city” areas and, further, that highly victimised young people in these areas are under-represented in the BCS.

To examine the assertion that the problem of victimisation for young people is particularly acute in “inner city” areas (as defined by the BCS), we examine the prevalence of violence, by age group, for areas defined as inner city as against all other areas. These areas make up just over ten per cent of the BCS sample. It is worth emphasising that this definition, corresponding broadly to ‘metropolitan’ cities, includes many affluent areas as well as areas of relative poverty and deprivation. This analysis is shown in Figure 3.4.

**Figure 3.4 All violence - prevalence by age within inner city**




For young people aged 16-24, the prevalence of violence is very similar in inner city and non-inner city areas; in fact the prevalence for the second youngest age group (20-24) is slightly lower in inner city areas (although this difference is not statistically significant). In the middle age groups (30-34 through 45-49), prevalence in inner city areas is higher than in other areas but the gap narrows in the older age groups, until ages 70-74 where it falls close to zero.

Alternative breakdowns are shown in Figures 3.5 and 3.6. In Figure 3.5, prevalence of all violence for respondents in the 10% most deprived areas (in England) is plotted against all other areas (in England). This time there is a slightly clearer difference between deprived and less deprived areas, with a higher prevalence of victimisation in deprived areas across the age range. However, as in Figure 3.4, the difference is much smaller in the youngest and oldest age groups.


Figure 3.6 shows prevalence of all violence for the ACORN category 'Hard pressed' (just under 20 per cent of the BCS sample), the category in which there is the highest overall rate of violence. Only with this classification do we see a clear difference between area types for young people as well as older respondents, with the largest difference (nearly 4 percentage points) being for those aged 16-19.

If we were to begin with the assumption that young people in deprived/inner city areas were much more likely to be victims of violence than those in other areas, then the above analyses could be interpreted as *suggestive* of bias in the BCS. On other hand, if we are prepared to accept the above analyses as representative, on the basis that they are the best evidence that we have, then it is difficult to conclude with any degree of certainty that young people in such areas are disproportionately affected by violent crime, as compared with young people in other areas. Whatever the case, it is clear that, a young person does not need to live in a deprived/inner city area to have a relatively high probability of being a victim of some sort of violence within the BCS reference period.

**Figure 3.5 All violence - prevalence by age within 10% most deprived areas (England only)**



**Figure 3.6 All violence - prevalence by age within 'Hard pressed' ACORN category**




## Incidence of violence by age and area type


One possible scenario is that, while prevalence for young people is not disproportionately higher in deprived/inner city areas, young people that are victimised in these areas are more likely to be repeat victims than their counterparts in other areas. To test this idea, we examine incidence rates (as opposed to prevalence); these are shown in Figures 3.7 and 3.8 (which are equivalent to Figures 3.5 and 3.6). Incidence rates are expressed as the number of incidents reported over the 5-year period under examination.

The incidence rates display very similar patterns to the prevalence rates, both in the rate of decrease across the age range and the differences between deprived/high crime areas and other areas by age group. However, this suggests that, amongst BCS respondents, there is no greater tendency for victims of violence (of any age) in deprived/'high crime' areas, to be repeat victims of this type of crime, as compared with victims in other areas.

**Figure 3.7 All violence - incidence by age within 10% most deprived areas (England only)**



**Figure 3.8 All violence – incidence by age within 'Hard Pressed' ACORN category**



Together these charts suggest that, whilst being young makes an individual much more likely to be a victim of violence and living in an deprived/high crime area may compound the risk of victimisation to a small extent, there is no *additional* risk from the combination of being young and living in a deprived/high crime area. Put more formally, there would appear to be little or no interaction between age and living in this type of area with respect to the risk of being a victim of violence.

## Assumptions

Of course, this interim conclusion depends upon the assumption that the achieved sample for the BCS is just as representative of young people in deprived/inner city areas as it is of young people in other areas. This assumption is certainly open to question but is rather difficult to test.

We might re-formulate the assumption thus: in deprived/inner city areas young people who do not respond are no more different from those who do respond than young people in other areas. We cannot test this assumption directly because, as is the case with all surveys of this nature (i.e. those using face-to-face interviewing at addresses randomly sampled from the PAF), we have little or no information on those young people who do not respond to the survey.

(Although we can identify cases where a person was selected but refused or was unable to take part, we cannot identify, using the available data, the age or sex of that selected person<sup>15</sup>. In other cases, e.g. where the interviewer failed to make any contact or the person that answered the door refused on behalf of the household, there would be no possible way of knowing if a young person would have been selected.)

While we cannot test this particular assumption, there still exists the possibility that differential response amongst young people, whether related to age or area type, could in fact be the cause of “hidden” bias in the survey estimates. For example, prevalence of violence may not be constant within age group. Whilst the data are calibrated to counter the effects of differential response rates by age, gender and region, adjustments for age are made by age group (16-19, 20-24, 25-34 and so on) as opposed to individual age. Therefore, if response amongst 20-21 year olds, for example, was significantly lower than that amongst 23-24 year olds, some downward bias might still exist in the calibrated results.

In the next section we investigate the potential for this kind of bias but first we look at response rates by grouped age. This will illustrate the extent of the adjustment made by the calibration weighting.

### 3.4 Response rates by age

We cannot calculate response rates by age directly as we do not have sufficient information on non-respondents. However, we can compare the composition of the achieved sample with population estimates. In fact, this sort of analysis is done each year in the BCS technical report, using age-sex categories (men/women aged 16-19, 20-24, 25-34, 35-44 etc). This analysis shows that the core sample (weighted by the appropriate sample design weights, but not taking account of the further calibration weighting carried out by Home Office researchers) typically under-represents those aged under 35.

For example, in 2007-08 this analysis showed that men aged 16-19 made up 5.7% of the weighted core sample and 6.9% of the population of men; those aged 20-24 made up 7.0% of the sample and 8.8% of the population; and those aged 25-34 made up 14.8% of the sample and 16.4% of the population. Slightly smaller differences were found for women in the corresponding age groups but in each case the core sample was shown to under-represent both men and women in these age groups. The same analysis was published in the previous year (2006-07) and the following year (2008-09) and very similar patterns were shown.

---

<sup>15</sup> Fieldwork documents show that interviewers are asked record the sex and broad age where an individual has been selected. However, this information was not available to NatCen.


We have taken this approach a step further using the same five years of survey data (2003/04 to 2007/08) as used in all the above analyses to estimate response rates for each age-sex category. We have compared the age/sex composition of the combined achieved sample (weighted by the appropriate sample design weights) with mid-year population estimates for 2005. To obtain the estimated response rates we have simply taken the sample percentage, divided by the percentage in the population and multiplied by the overall response rate (74.4%). The results (for the five years combined) are shown in Figure 3.9. This shows some interesting patterns.

(It should be noted here that application of the appropriate design weights is necessary to make valid comparisons between the achieved sample and the population. This is especially true if we are interested in the relative response propensity by age, as young people aged 16-29 - particularly those aged 16-18, most of whom live with their parents - are much more likely to live in large households and therefore have a much lower probability of selection for the survey. Even if the survey response rate were 100%, the un-weighted sample would under-represent young people, as well as those PFAs which were under-sampled. Therefore, adjustment for differential selection probabilities, by PFA, dwelling unit and household size, is necessary to estimate accurately individual-level non-response.)

The lowest response rate of all (44%) is for women aged 85+ (1.7% of the population). For those under 85, the lowest response rates are found in the 20-24 age group; this is the case for both men and women. After this response rates for both genders climb gradually to a peak and then drop off in later years. For women, response rates peak at 86% for those aged 45-54; up to this point response rates are higher than those for men in each age group. After this, however, response rates for women start to fall, whilst for men they continue to climb, overtaking those for women in the 55-64 age group and peaking at 86% in the 65-74 age group. For those aged 75 and above, response rates for men are appreciably higher than those for women; however, given the very low levels of victimisation in these age groups this is of academic interest only.




**Figure 3.9** Estimated response rates by age within sex




## Response rates for young people

Next we examine estimated response rates for individual ages from 16 through to 29 (within sex), based on our five years of combined data. The results are shown in Figures 3.10 and 3.11, plotted against average incidence of violence for individual years of age over the same five year period. This allows us to see whether differential response within age group could be a source of “hidden” bias.

**Figure 3.10** Estimated response rates by individual age vs. incidence of violence (males)



**Figure 3.11 Estimated response rates by individual age vs. incidence of violence (females)**



Looking first at Figure 3.10 (analysis for males), response rates for men are well above 60% for 16-18 year olds and then fall abruptly to a low of 55-56% for those aged 19-21<sup>16</sup>. After this initial fall, response rates continue to fluctuate but appear to be on an upward trend by the late 20's.

In terms of the potential for bias, the key group to look at is 20-24's where response rates are higher in the latter ages whilst levels of violence are falling quite rapidly (ignoring a "blip" at age 23, probably caused by nothing more than random "noise" in data) across the age group.

Turning to Figure 3.11 (females), we see a similar pattern as for males with those aged 16-18 achieving higher response rates than those aged 19-21. Thereafter, female response rates increase gradually (with minor fluctuations) to around 70% for the late 20's.

For female respondents to the BCS, while the underlying trend in response rates from age 19 is a gradual increase with ascending age, the trend in levels of violence is a gradual decrease year on year. As with males, (particularly males aged 20-24 as seen above), this combination will inevitably result in a small amount of bias in the overall survey results. However, given the relatively small differences in response rates, this bias is almost certainly non-significant (i.e. within confidence limits).

Whilst it would be interesting to know whether these patterns are apparent in individual survey years, it would not be possible to draw any firm conclusions from repeating this analysis by year due to the small numbers involved. Combining five years of data has enabled us to estimate response rates with reasonable precision. However, given that there are fewer than 200 respondents of each age in many years, it is not possible to replicate this analysis for individual years of data. Examination of the raw numbers (weighted by the design weights), however, shows that there is a consistent tendency for lower response amongst those aged 19-21 as compared with those aged 16-18: this pattern is apparent in each of the five years (2003/04 to 2007/08).

<sup>16</sup> Most 16-18s are resident in the household of their parents; at age 19-21 the influence of parents on response appears to be reduced.

## Response rates for young people in deprived/inner city areas


Unfortunately we cannot reproduce this analysis for deprived/inner city areas, the problem being that detailed population estimates for such areas are not available. However, ONS do publish population estimates for Lower Super Output Areas by broad age and sex as experimental statistics and we have been able to use these to create (mid-2005) population estimates for areas in the lowest decile of IMD. As in the above analysis, we have compared the population composition with the achieved sample (weighted by the appropriate design weights) in order to estimate response rates, using our combined survey data (2003/04 to 2007/08). The results are shown below in Figures 3.12 and 3.13.

For young people aged 16-29, response rates in the 10% most deprived areas are markedly lower for males at 54% as against 62% in other areas. However, there is only a small difference for females (66% versus 65%). In older age groups, levels of response are lower in deprived areas with larger differences for males than females.


Looking at the corresponding incidence of violence for those aged 16-29, young male respondents in deprived areas appear to be no more likely to be victimised than those in other areas. Young female respondents, on the other hand, experience over 60% more violence if they live in deprived areas compared with those in less deprived areas (although they are still less likely to be victims than males in the same age group). This difference is also evident for older females aged 30-44 and 45-59.

This is an interesting finding and suggests that, contrary to expectation perhaps, it is response rates amongst young females in deprived areas/inner city areas rather than those for young males, which should concern us. The good news is that, according to this analysis, the BCS achieves similar levels of response for young females in deprived areas as it does for young females in other areas.

**Figure 3.12** Estimated response rates by broad age for 10% most deprived areas versus other areas (males)



**Figure 3.13** Estimated response rates by broad age for 10% most deprived areas versus other areas (females)



We now turn our attention to an examination of response rates at total sample level, with particular focus on area-level variables.

### 3.5 Response rates by area type (and interviewer observed variables)

Table D1 (see Appendix D) shows response rates for all relevant variables in the issued sample files from 2002-03 onwards. The variables fall broadly into two categories, area-level information and interviewer observations.

Response rates were calculated using the following formula:

$$(\text{Full interviews} + \text{Partial interviews}) / (\text{Total addresses} - \text{Ineligible addresses})$$

They were calculated for each tranche of issued sample and for all tranches combined.

Because the balance of the sample by PFA has changed over the period, the data were weighted by the PFA selection weights. (Comparison with the raw response rates shows that this step mostly had a negligible impact but it improves the comparability of results across years.)

The main purpose of the breakdown by year is to examine the grounds for the expectation that any patterns we see overall are also observable within tranches and that they remain reasonably consistent over time. This is difficult to test formally but detailed inspection reveals that where response rates were low, they were consistently low and vice-versa. For example, the relative response rates across ACORN 2001 groups are remarkably consistent year on year: The lowest response rates were always found in the groups Educated Urbanites and Inner City Adversity, whilst Wealthy Executives and Affluent Greys have consistently experienced the highest response rates of the 17 groups. Another notable example is for ONS Ward classification groups where Metropolitan professionals, Deprived city areas and Inner city estates had the three lowest response rates of the 14 groups.

When looking at the combined data (from 2002-03 onwards), all variables apart from two (Rotated or fresh PSU and location on main road) were found to be significantly related to response.

## Categories with low response rates

When looking at those variables that are related to response, the following categories stand out as having consistently low response rates compared to the overall average (74%):

- *Region/Area type*: London (62%); Inner city areas (68%)
- *From ONS Ward classification*: Metropolitan professionals (62%); Deprived city areas (59%); Inner city estates (57%)
- *From ONS Ward classification 2001*: Multicultural Areas (68%); Inner City Multicultural (58%); Prospering Metropolitan (57%)
- *From ONS District Classification (Family)*: Education Centres and Outer London (66%); Inner London (52%)
- *From Acorn Group*: White Collar Workers (69%); Affluent Urbanites - town & city (67%); Prosperous Professionals - met areas (61%); Better Off Executives - inner city (60%); Council Estates, high unemployment (60%); Multi Ethnic low income (59%)
- *From Acorn Group 2001*: Aspiring Singles (67%); Asian Communities (67%); High Rise Hardship (64%); Educated Urbanites (58%); Inner City Adversity (58%)
- *From Crime and Disorder Index*: 10% most deprived wards
- *Visible*: Security gate (70%); Estate/block security lodge/guards (63%); Entry phone (61%)
- *Location of dwelling*: Above shops (64%)
- *Type of accommodation*: purpose built flat (64%); converted flat (63%); rooms, bed-sitter (68%)
- *Type of entrance*: Common entrance: lockable (61%); Common entrance: not lockable (67%)
- *Physical condition of property*: Fairly bad (69%); Very bad (66%)

Many of these types of areas/addresses are ones in which lower than average response rates are observed on many household surveys. However, it is interesting to note, in the case of the BCS, both:

i) The level of (non-)response in some areas, for example 'Inner city estates' and 'Deprived city areas' from ONS Ward classification have particularly poor response rates (below 60%).

ii) The inclusion of some more affluent areas e.g. ACORN groups 'Affluent Urbanites - town & city' and 'Prosperous Professionals - met areas', illustrating that non-response is not solely concentrated in

deprived inner cities. For example, high-status accommodation, with systems such as entry-phones to restrict access, is also associated with lower survey response rates.

## Comparison with other surveys

To examine the extent to the patterns we have seen above are typical of large-scale face-to-face social surveys, we have attempted to compare area-level response rates on the BCS with other surveys conducted by NatCen, namely the Health Survey for England (HSE) and the British Social Attitudes Survey (BSAS). The results for 2007-08 are shown in Table 3.1 below and for comparison are restricted to England only.

Table 3.1 Comparison of BCS response rates with other social surveys (2007-08) (England only)			
	BCS	HSE	BSAS
Overall response rate	75%	66%	53%
London	64%	56%	45%
Lowest IMD quintile	71%	65%	50%
<i>ONS Ward classification 2001:</i>			
Multicultural Areas	69%	62%	44%
Inner City Multicultural	60%	56%	46%
Prospering Metropolitan	59%	57%	44%

In each case, response rates in London were appreciably lower than other regions; this is typical of all household surveys. Similarly, addresses classified as Multicultural Areas, Inner City Multicultural and Prospering Metropolitan had amongst the lowest response rates on all three surveys, but the gap between these and the overall response rate was larger for the BCS. The 20% most deprived areas also had lower than average response rates; in each case there was only a few percentage points difference from the overall rate but again the BCS showed the largest difference.

On the other hand, BCS has a much higher level of response, both overall and in all the areas shown, so whilst some areas show relatively poor levels of response on the BCS, the actual response rates for these areas are appreciably better than those found on other surveys.

## Modelling non-response

These response rates give us a flavour of the broad patterns of non-response. However, many of these variables are correlated. For example, inner city areas are likely to have higher concentrations of flats and properties in poor physical condition.

In order to understand which of these factors are important and which are by-products of other (related) factors, the next logical step might be to estimate multivariate model(s) of non-response. The ultimate aim would be to find out whether adjusting for non-response, as part of the weighting scheme, would have any significant impact on the survey estimates.

We might ideally like to be able to re-build the weights for a given year by including further non-response adjustments prior to calibration. This is not as easy as it sounds, however, not least because it would be a complex task to re-run (and correctly replicate) the final calibration step.

### 3.6 Effectiveness of calibration weighting

An alternative way of answering the question is to examine the extent to which the calibration that is currently carried out, adjusts for some of these aspects of non-response indirectly. For example, calibration by region should adjust for some of the area-level variation in response.

We can estimate the profile of the population (for all the variables in the issued sample files) by taking the eligible sample (i.e. removing the ineligible addresses) and weighting it by the PFA selection weights. By comparing this profile with the profile of the achieved sample weighted by the final (calibrated) household weight, we can see the extent to which the calibration weighting adjusts indirectly for various aspects of household non-response.

The analysis described was carried out for all survey years from 2002-03 onwards and is shown in tables B2a-B2c. The comparison is made by calculating the ratio between the proportion of each category in the eligible sample and the proportion in the (final) achieved sample. Any ratios greater than 1.1 (i.e. where the proportion in the achieved sample is more than 10% higher) or less than 0.9 (i.e. where the proportion is less than 10% lower in the achieved sample) are highlighted.

Notable categories that tend to remain under-represented in the final achieved sample, include:

- *From ONS District Level Classification: Group (2003):* London Centre
- *From ACORN Group (2001):* High Rise Hardship; Inner City Adversity
- *Visible:* Entry-phone
- *Accommodation type:* purpose built flat
- *Condition of house/flat:* Fairly/Very bad

Each of these categories, are areas/addresses with higher than average crime rates on the BCS. However, none of these categories is under-represented to a very large extent and, moreover, they are not very substantial sub-groups in the whole BCS sample.

The most striking aspect of this analysis is the extent to which the calibration weighting brings the achieved sample into line with the “correct” profile, as evidenced by the relatively small number of shaded cells in Appendix Tables D2a-D2c. This indicates that to a large extent, the calibration step adjusts for many aspects of non-response, in just the way it is intended to do. This suggests, in turn, that additional non-response adjustments, in the form of a non-response model for example, would have little impact on the profile of the achieved sample and hence little impact on the weighted results.

### 3.7 Modelling victimisation (violence)

All our analysis so far points to the conclusion that there is only a relatively small degree of bias in the survey results, unless those not responding to the BCS are significantly different from the respondents in ways we have not measured (and cannot estimate), for example if the young people who were not interviewed were substantially more highly victimised than those who were interviewed. (As we have seen, this small degree of bias would appear to stem from differential response/victimisation within the youngest age groups and the slight under-representation of certain types of area and property as listed above.)

The question of whether or not non-respondents are more highly victimised is a difficult one to answer satisfactorily but the small amount of available evidence points to the conclusion that they are not. In 1996, NatCen attempted to ask those selected who refused to take part whether they had been victims of some key types of incident. The conclusion of this (not entirely successful) exercise was, if anything, those who refuse to take part in the BCS are in fact less likely to be victims than respondents to the survey<sup>17</sup>. This is consistent with the findings of a study of the Scottish Crime Survey (albeit in a period in which data collection was carried out by telephone) which concluded that refusal seemed to be the principal cause of bias in favour of victims<sup>18</sup>.

However, we cannot discount the possibility that non-respondents are, in fact, somewhat more highly victimised than respondents, particularly in the youngest age groups, where the analysis has demonstrated that response rates are lower than for the BCS as a whole. To get a feel for whether this might be possible, we have attempted to model victimisation using logistic regression. All relevant variables available from the issued sample file and the survey data files were included. In line with other analyses in this section, we have done this using a combined dataset of 5 years of survey data (2003/04 to 2007/08).

The 20 predictors that went into the model fall broadly into three categories: individual/ household characteristics, lifestyle indicators and area-level information. The variables are listed below:

*Individual/ household characteristics:*

- Age/sex
- Marital status
- Ethnicity
- Tenure
- Highest educational qualification
- Employment status
- Total household income (6 categories)
- Household structure
- Disability/illness

*Lifestyle indicators:*

- Hours spent away from home during the day
- No of visits to pub/wine bar last month
- No of visits night club bar last month
- Vehicle ownership

*Area-level information:*

- Inner city
- GOR
- ACORN group 2001
- How common is vandalism

---

<sup>17</sup> Lynn, P (1997) 'Collecting data on non-respondents to the British Crime Survey. Unpublished report to the Home Office'. London: Social and Community Planning Research.

<sup>18</sup> Hope, S (2005) 'Scottish Crime and Victimisation Survey: Calibration Exercise: A Comparison of Survey Methodologies'. Scottish Executive. [www.scotland.gov.uk/Publications/2005](http://www.scotland.gov.uk/Publications/2005).



- How common is rubbish
- How common are homes in poor condition

*Other:*

- Year of survey data

The dependent variable in the model indicated whether or not the respondent had been a victim of violence. The model was weighted by the design weights and the variables age/sex and GOR were forced into the model first before other variables were allowed to enter on a stepwise basis.

The final model is shown in Appendix Table D3. Only four variables do not appear in the model, these are:

- Total household income (6 categories)
- Inner city
- How common is vandalism
- How common is rubbish

The Wald statistic in the final model suggests that age/sex is the strongest predictor of victimisation, even after controlling for all the other covariates in the model. The model confirms the aggregate-level analysis that showed that males aged 16-24 were most at risk from violent crime.

## Characteristics associated with increased risk of violent victimisation

Other notable characteristics associated with an increased risk of violent victimisation are as follows:

- Being divorced or separated (compared with being married)
- Being White (compared with being Asian or Black)
- Living in rented accommodation (compared with owning the accommodation)
- Being unemployed (compared with being employed)
- Being a lone parent (compared with living in a household with no children)
- Having a long-term illness, particularly one that is limiting in day to day activities (compared with not having an illness)
- Visiting a pub at least once in the last month (compared with not having visited a pub)
- Visiting a nightclub at least once in the last month (compared with not having visited a nightclub)
- Not owning a vehicle (compared with owning a vehicle)
- Residing in areas classified as ACORN Groups such as 'High Rise Hardship', 'Asian communities' and 'Aspiring Singles' (compared with 'Wealthy Executives')

Perhaps the most interesting aspect of the model is the lifestyle characteristics. Visiting pubs or nightclubs three or more times a week raises the odds of victimisation by a factor of at least 1.8 (for pubs) and a factor of 1.8 (for clubs). This means that someone who visits both pubs and clubs on a regular basis is just over 3 times more likely to be a victim of violence than someone who does not frequent such establishments. When combined with the increased odds associated with being young and male, this creates a substantial risk of violent victimisation. (Moreover, there may be other aspects of lifestyle not measured by the survey that are also related to the likelihood of victimisation.)

We have seen that young men (and to a lesser extent young women) have significantly lower than average response rates. If lifestyle characteristics such as visiting pubs/clubs are also related to the probability of response (which would seem plausible as these people may spend more time out of their homes) this raises the possibility that non-respondents (particularly in younger age groups) might be more highly victimised than their responding counterparts.

Given that we do not have accurate information on how often young men and women in general visit pubs/clubs, we cannot gauge the extent to which the achieved sample under-represents those that frequent them. However, if we make some assumptions about the degree to which non-respondents *may* experience greater victimisation, we can re-calculate the survey estimates and compare these with actual estimates for the same period.

### 3.8 'What if' analysis

This was done for violence using our five years of combined survey data (2003/04 to 2007/08).

First we calculated incidence rates for respondents for five-year age/sex groups (i.e. males/females 16-19, 20-24, 25-29 etc). We then calculated incidence rates for non-respondents, based on an assumption about their relative degree of victimisation (compared to respondents in the same group). To do this we simply multiplied the incidence rates for respondents by a constant 'inflation' factor. (For example, if we wished to assume that non-respondents experienced 10 per cent more victimisation incidents than respondents we used a factor of 1.1; 20 per cent more required a factor of 1.2 etc).

By combining the two sets of estimates (for respondents and non-respondents) taking into account the response rate in each group, we derived estimates for the number of violent incidents in each age/sex group. These were then grossed up to population numbers to give estimates of the total number of incidents in each age/sex group in the population.

Summing these figures across all the age/sex groups, dividing by the population total and multiplying by 10,000 gave us our 'new' (annual) estimate of violent incidents per 10,000. This was then compared with the overall rate for respondents *plus* a typical confidence interval for the incidence of violence.

Over the five-year period the average incidence was around 570 per 10,000, whilst a typical 95% confidence interval was around +/- 50 giving a 'typical' upper limit of around 620. By changing the inflation factor we were able to gauge the degree to which non-respondents would have to more highly victimised to produce a 'new' estimate greater than the upper limit of the confidence interval.

An inflation factor of 1.25 was enough to produce an estimate that falls (just) outside the 95% confidence interval. In other words, we estimate that non-respondents (of all ages, both male and female) would have to experience at least 25% more victimisation than respondents for there to be a (statistically) significant downward bias in the annual estimates of the incidence of violence.

In 1996 survey<sup>19</sup>, the incidence of violence was estimated to be 1,046 per 10,000. Using the same model as above, an inflation factor of just over 3.5 was required to produce an estimate of this magnitude.

---

<sup>19</sup> BCS estimates for 1996 represent the 'peak' victimisation rates over the whole period of the BCS since 1982, whilst the survey also experienced the highest response rate over the period (83%).

## 3.9 Summary and conclusions

### Summary

We have seen that young people are disproportionately affected by personal crime and, in particular, violent crime; moreover, living in a deprived/high crime area further increases the likelihood of victimisation, particularly for women. However, we found no evidence to suggest that there is any *additional* risk from the combination of being young and living in a deprived/high crime area over and above the risks associated with each factor alone.

We have shown that response rates for young people, in particular young men, are much lower than average and for the some age groups are (effectively) as low as 55% (young men aged 19-21). When we estimated response rates for young people aged 16-29 in deprived areas, we found that, compared to other areas, they were lower for males but not for females. So whilst young females in deprived areas experienced over 60% more crime than their counterparts in other areas, response rates for the two groups were similar.

We looked at area-level response rates and found that response rates in some areas, for example 'Inner city estates' and 'Deprived city areas' from ONS Ward classification, were relatively low. These patterns were found to be similar year on year. However, when compared to other surveys such as the HSE or BSA, response rates on the BCS were generally higher, including in deprived/inner city areas.

Despite some relatively low response rates, calibration weighting does an effective job of adjusting (indirectly) for non-response, both at area-level and by type of property. This, in turn, suggests that, unless non-respondents are very different from respondents, there is not likely to be any significant bias in the current survey estimates.

We have estimated a statistical model of the likelihood of violent victimisation. This shows that, in addition to being young, male and living in high crime areas, other factors such as having a long-term illness, being unemployed and frequenting pub/clubs increase the overall risk of experiencing violent crime.

Whilst we are unable to use this information directly to estimate whether those not responding to the BCS are more highly victimised than those who do respond, we have been able to calculate that non-respondents would have to experience at least 25% more victimisation than respondents for there to be a (statistically) significant downward bias in the annual estimates of the incidence of violence. Moreover, in order for victimisation rates to reach mid-90's levels, non-respondents would have to experience 250% more victimisation than respondents.

### Conclusions

In conclusion, we have found no compelling evidence to suggest that the survey results are significantly biased due to the under-representation of young people in certain types of area. Moreover, the analyses of response rates and victimisation rates (prevalence and incidence) strongly suggest that if any bias does occur, then it is very likely to be consistent year on year. In that case, such bias would not affect the ability of the BCS to identify the overall trends in victimisation within the population of England and Wales.

## 4 The period covered in annual reports

### 4.1 Background and task

This chapter considers an alternative approach to the analysis and reporting of experience of crime to that currently provided in Home Office publications. Reflecting the continuous fieldwork design of the BCS, the prevailing measure reports on crimes experienced in the year before their interview by individuals *interviewed* in a particular year. It does *not* provide an estimate of crimes experienced in a *specific* year – instead the measure covers a period of 23 months. The aim of this chapter was therefore:

- To assess whether and how a more intuitive measure could be constructed to provide the incidence of crimes *occurring* in a particular year;
- To investigate the *trends* that this measure provides over time and compare them to those observed in the prevailing measure.

Prior to 2001, BCS fieldwork was carried out in the early part of a calendar year and asked respondents about their experiences of crime in the preceding calendar year. Published statistics focused on offences committed in that calendar year.

In 2001, the BCS adopted a continuous fieldwork design that was representative of the population resident in households in England and Wales each quarter. The reference period for experience of crime was the full 12 month period prior to the month of interview (so for an interview in April 2009 the reference period for analysis would be April 2008 to March 2009). This change from a reference to the previous calendar year was required due to the problems for recall that would be presented for interviews later in the year.<sup>20</sup>

Since the introduction of the new design reporting has been based on *interviews* carried out in a particular year (primarily the year April to March, but also rolling years; i.e. January to December, July to June etc). Rather than a national measure of the incidence of victimisation in a 12 month period, this provides a measure over a 23 month period, with the highest weight in the estimate being given to the months in the centre of the distribution. This is illustrated in Figure 4.1. For example, for the year of interviews from April 2007 to March 2008 the reference period for crimes includes the months from April 2006 to February 2008 (as indicated by the diagonally shaded area) with March 2007 the central month. Assuming 50,000 interviews are achieved each year, for an estimate of crime based on interviews in 2007/8 April 2006 would be within the reference period for 4,167 respondents compared with 50,000 who would include March 2007.

For an estimate of crime *incidents* falling within that year (April 2007 to March 2008) the centre of the reference distribution would be September and October 2007 – or around six months later than for the measure based on a year of interviews (and clearly only a 12 month period rather than one of 23 months).


---

<sup>20</sup> See the methodological review carried out by Lynn and Elliot that recommended these changes at <http://www.homeoffice.gov.uk/rds/pdfs08/bcs-methodology-review-2000.pdf>

Figure 4.1 also shows the interviews that would contribute to a measure of crime incidents occurring in 2007/8 (green shading). The reference period for people interviewed in each of the 23 months from May 2007 to March 2009 would cover at least one month of 2007/8. For those interviewed in April 2008, their reference period would cover the full 12 months April 2007 to March 2008, whereas for those interviewed in March 2009 their reference period would only cover the final month – March 2008. Looking at the total number of cases whose reference period covers each month of 2007/8 in the example, there are 50,000 in each month.

Note that to look at part of a reference period requires us to move away from an annual measure of crime. Once it is accepted that we are not looking at a full year reference period and that we are looking at monthly rates, there would be no reason not to use a smaller number of interviews to look at the crimes.

Figure 4.1 Example of interview months and crime reference periods for interviews between April 2007 and March 2009



## 4.2 Calculating annual incidence rates

### A rate per month approach

Home Office publications of BCS findings report 'incidence rates per 10,000 individuals' (or households, depending on the type of crime). For any particular crime, the incidence rate for a given period is the mean number of crimes experienced per person (or household) multiplied by 10,000 to make the measure easier to interpret. As such, it includes multiple crimes against single individuals. It is this latter point that distinguishes an 'incidence rate' from 'prevalence rate' in Home Office reports – the latter is the proportion of individuals (or households) who are the victims of a particular type of crime one or more times.

The approach adopted for calculating incidence rates on an annual basis was one of first calculating an incidence rate per month (i.e. an average number of crimes per individual or household during the month). Months define the reference period and respondents are asked in which month incidents happened during the interview. Rates can be calculated for any given period, and their arithmetic property allows us simply to add estimates together to obtain rates for longer periods. Crucially, we can use different (and overlapping) samples for each of the months in a particular period and still combine them as estimates of a specific period.

The first step was to combine data from 2001 to 2009 interviews and resolve minor differences in the relevant variables between years. With this in place, the process of validating incidents reported to the survey (in terms of their nature and timing) that was carried out for the published interview year rates was replicated. In addition, the process of capping of series incidents was replicated (a maximum of five incidents are recorded in a series to reduce the impact of outliers on estimates). In this way, the same number of incidents across the full sample was identified for the analysis.

### Assigning incidents to months

A requirement of the monthly rates approach is to assign incidents to specific months. Whilst respondents are asked in which month an incident occurred (in order to ensure the incident falls within the 12 month reference period), there will be a degree of uncertainty in some instances. In these cases, a form of imputation is required to assign incidents to months on the available information. We need to be confident that the impact of misallocation of incidents to months is minimal to be able to rely on comparisons with year of interview measures.

To state the principles of the assignment of incidents to specific months:

- Where respondents were able to give a specific month for a single incident, this was the month to which the incident was assigned;
- Where they could not provide a specific month but identified a quarter, the incident was assigned to one of the three months in that quarter at random;
- Where they were not able to assign the incident to a quarter, the incident was allocated at random to one of the 12 months in the reference period;
- Where there were series of incidents, for instances where the number of incidents was at or below the capping limit of five and where it was known in which quarter each incident took place, each incident in the series was randomly assigned within the possible three months;

- Where there was information for some and not for others, incidents which were not already accounted for were allocated to quarters where there had been uncertainty (and randomly assigned to a specific month subsequently);
- Where there was no information about the quarter, incidents were assigned randomly across the 12 months;
- For series incidents which had been capped to five, the first step was to cap the number in any one quarter to five, and then reduce the number across the quarters at random until there were five incidents in the year.

Table 4.1 provides a summary of how incidents were allocated to months for all incidents in the surveys between 2001/2 and 2008/9 (there was little variation between years). For nearly two-thirds of cases (63 per cent) it was possible to simply use the specific month that the respondent gave for the incident.

In a further 36 per cent of cases it was possible to assign incidents randomly within a known quarter. This largely applied to series incidents, for which respondents are not asked to provide a specific month. Seven per cent of incidents were part of a series where the number was capped at five incidents. In these instances, the quarterly totals were reduced iteratively in a random fashion until a maximum of five incidents remained for the year.

The random month in a year was assigned in less than one per cent of cases. It might be argued that a different approach would be to take account of the seasonality of crime rather than use purely random allocation, but the numbers here are very small. In terms of the random allocation to months within a quarter, trying to take account of seasonality would be unlikely to improve estimates. Nevertheless, a small improvement for future analysis might be to improve the basis on which incidents are allocated, particularly where there is random allocation across a full year.



**Table 4.1 Summary of how incidents were assigned to months (all incidents 2001/2 to 2008/9)**

	% (all incidents)	Number of incidents
<b>Allocation to specific month by respondent</b>	<b>63.5%</b>	<b>91,400</b>
<b>Random allocation of incidents within a known quarter (net)</b>	<b>36.0%</b>	<b>51,865</b>
Single incidents where only the quarter was known	0.4%	511
Series incidents where the quarter was known for all incidents and total not capped	28.2%	40,594
Series incidents where the quarter was known for all incidents but quarterly totals capped at random to maximum of five across all quarters	7.5%	10,760
<b>Random allocation of incidents within a year (net)</b>	<b>0.5%</b>	<b>711</b>
Single incidents where only year known	0.1%	118
Series incidents where quarter not know for one or more incidents in the series (random allocation of remaining incidents between quarters where incidents known to have occurred)	0.4%	593
<b>Total</b>	<b>100.0%</b>	<b>143,976</b>

It should be noted that the random allocation of incidents to months will mean that the estimates produced will vary slightly each time they are recalculated. This may become an issue were reporting of annual rates to become routine.

## Calculating annual estimates from the monthly rates

Specific offence categories were created in line with the approach for the year of interview rates. Having assigned incidents to specific months, it was then possible to calculate means for each of the 107 months in the period April 2000 to February 2009 for each of the offence categories. These means could then be held in a database and rates for any period of months constructed. For the annual measure, the estimate is simply the sum of the rates for the 12 individual months. This measure is referred to in the remainder of this report as the 'annual incidence rate', and is compared to the 'year of interviews' measure currently in use by the Home Office.

## Weighting issues

Weighting the new annual incidence rate measures presents some significant challenges principally because sample cases are used different numbers of times depending on how many of their reference months are within the period of interest.

Calibrated weights are calculated separately for each year of interviews by the Home Office. The weights used in this analysis are simply those provided with each year of data - 2001/2 etc).<sup>21</sup> The

<sup>21</sup> The mean values of the calibrated weights varied widely between the datasets due to changes in the process at the Home Office between years. In order to facilitate analysis of cases across different years, the weights were rescaled within each year to a mean value of one. The rescaling does not affect relationships between cases

application of these weights during the calculation of monthly estimates addresses the major issues of differences in sampling probabilities and response. However, an approach from first principles might be to recalculate the calibration weights for the *particular set* of cases included in each monthly rates calculation with reference to population estimates for month of interest. A consideration here is that this would be a significant undertaking and may add little to the accuracy of estimates beyond simply applying the weights provided with each financial year dataset.

A further issue relates to the nature of interviewer assignments for the survey. The BCS is designed to be balanced by quarter in terms of Primary Sampling Units and fieldwork assignments, and design weights are also designed to balance each quarter. A sample in an individual month may not be representative in the same way that the quarter as a whole will be, depending in particular on the nature of management of assignments during the particular quarter (i.e. whether they are largely issued at the start of the quarter or at intervals during the quarter). This is an issue in that for two out of every three monthly rates estimates we use interviews from parts of two quarters (covering the beginning and end of the reference period). However, this is only likely to have a small impact on the estimates for a particular month as three full quarters would be used. Furthermore, the months used would be, for instance, the last month from one quarter and the first two from another and so would reflect the variation in patterns of assignment across quarters.

## Significance testing

Testing the statistical significance of differences between annual incidence estimates provides further challenges. Any approach needs to take account of the fact that individuals will in some cases provide data for both of the annual estimates to be tested. The nature of the derivation of the estimates also involves the summing of monthly estimates outside a statistical analysis package.

A simple approach could approximate the variance of the difference between estimates for two years by  $2V(1-p\rho)$ , where  $V$  is the average estimated variance for each estimate,  $p$  is the proportion of individuals in common between the samples in the two years and  $\rho$  is the correlation coefficient between the relevant variable (total number of incidents in the year) for the two years. The standard error of the difference between the estimates can be derived by dividing the square root of this variance (to get the standard deviation) and dividing by the square root of the average sample size.

Clearly, this is a simplistic approach that does not take account of the elements of the complex sample design (clustering, stratification, weighting). Including this in a more sophisticated approach would by no means be an easy task and is beyond the scope of the current study. The problem would merit further work should it be decided that these were to become published trends.

## 4.3 How do the two measures react to changes in crime rates in a period?

Before looking at the trends that the new annual measure of incidence produces, consideration can be given to how, in principle, the derivation of the two measures defines how they will behave in relation to changes in crime levels over time. The features of the two measures give us three central points of difference:

- The annual crime rate measure is formed from crime occurring in a 12 month period, whereas the interview measure is affected by crimes occurring over 23 months;

---

within the sample. Without it, respondents from some years would have been over-represented in the estimates, which used cases from different years.

- Unlike with the annual measure, the interview measure is not affected equally by crime levels in these 23 months – months in the centre of the distribution have the biggest effect;
- The central point of the 23 months covered by ‘interviews in 2007/8’ is about six months behind the central point of the 12 months of ‘crimes in 2007/8’.


Perhaps the easiest way to understand the principles of the differences between the trends produced by the two measures is to simulate changes in crime levels and chart the results. Using fictional *monthly* rates of crime across several years, it is possible to derive the annual measure of crime (by summing the monthly rates for a period of interest) and to approximate the interview measure by taking the weighted average of monthly rates in a 23 month period and multiplying by 12 months (the weighting of the monthly rates in the calculation of the average relates to the proportion of respondents who would be referring to each month in the period – see the numbers of interviews listed in Figure 4.1 above). This is an idealised situation that does not arise in the comparison of actual trends – there is an underlying assumption that the survey reflects actual crime trends entirely accurately – but it serves to illustrate the behaviour of the measures.

Figure 4.2 displays how the measures react to rising and then falling crime rates in a four year period. The ‘annual’ line represents the new measure of the number of crimes that happened in a particular year (for instance the first point on the chart for this line is crimes in year Y1). The ‘interviews’ line represents the crime rate as measured by *interviews* in that *same* year (so the first point on the chart represents the crime rate for interviews conducted in year Y1). The ‘annual (lag 6 months)’ line is the annual measure, but simply moved to the right by six months, with the result that the middle of the reference period for this and the ‘interviews’ measure is around the same point on the chart (so the point for this measure at Y1 actually relates to crimes that occurred in the year starting six months before the start of Y1). Points on the chart relate to rolling annual periods one quarter on from the previous point – for example Y1+1Q is formed of the last nine months of the interviews / crimes contained in Y1 plus those from the subsequent three months.

The chart helps to illustrate several points:

- The trend for the annual measure clearly precedes that of the interviews. The peak of the annual measure is six months ahead of that for the interview measure, as we would expect with the 12 month compared to 23 month reference periods;
- For this reason, the annual measure with a lag of six months provides a more appropriate comparison (and closer trends);
- For a given crime rate over time, the result of the 23 month reference period will be a ‘smoother’ trend line than is seen with the shorter-reference period of the annual measure (although in practice, the two measures will be based on samples that are not entirely overlapping, and the sampling error between them will mean a smoother trend is not always observed);
- More specifically, because of the longer period, the interview measure will pick up a rise or fall in crime levels *before* the lagged annual rate;
- However, the rate of increase or decrease will be slower.

**Figure 4.2 Comparison of annual and interview measures using simulated data – rise and fall in crime levels**



Home Office reporting to date has been on separate annual periods rather than the rolling periods presented above. A close look at the separate annual periods within Figure 4.2 highlights an issue for the comparisons – that although the measures produce quite similar trends overall, the trends at some points are more divergent than at others. The point is highlighted in Figures 4.3 and 4.4. The charts relate to the same crime rate data and cover largely the same period, but we see that the trends are noticeably closer for the Y1 to Y4 series than for the Y1+2Q to Y4+Q series. This simply relates to the particular quarter in which changes occur or accelerate.

Figure 4.3 Comparison of annual and interview measures using simulated data – Y1 to Y4





Figure 4.4 Comparison of annual and interview measures using simulated data – Y1+2Q to Y4+2Q



The simulation here is of a neatly rising and falling crime rate, but there are other scenarios that would produce wider disparities between measures at particular points. For instance, a situation where there are sharp peaks at both ends of the 23 month reference period and a trough in the middle 12 months (where the annual measure would be focused) could produce a considerably higher rate for the interviews measure.

In addition to the divergence resulting from the difference in the lengths of the reference periods and the weighting of months within them, further differences will result from the nature of the calculation of the annual measure. One issue is that a (small) amount of variation will result from the random assignment of incidents to months within quarters and years in some instances. A second point is that the simple summing of monthly rates to provide the annual measure means that differences in the sample sizes for each estimate do not have a direct bearing on the annual measure – the same weight is given to a monthly rate estimate in a month whatever the sample size. This is a different treatment of cases compared with the interviews measure. These differences in the approach are likely to affect the closeness of low incidence crimes in particular.

## 4.4 The crime trends

The following charts display published incidence rates per 10,000 individuals (or households) for years of interviews alongside the newly developed annual rates for the major crime categories reported by the Home Office.

Although the major focus for discussion in Home Office reports is the measure of *numbers* of crimes, the focus is on incidence rates here as this allows a cleaner comparison to be made between the two series. The process for generating numbers of crimes is to take the incidence rate and multiply by an estimate of the size of the England and Wales population. Because different population estimates would be needed for the two series (due to the lag of six months between their central points – see below), this would add a layer of complexity to the interpretation of comparisons of the trends and was therefore not attempted.

Continuous fieldwork began in April 2001, providing a reference period going back to April 2000. The annual rates for the year 2000/1 is displayed but there is a relatively small sample size for the early months in that year, as there is for the later months in the final year 2008/9 (see Appendix Table E.2 for the sample sizes in each of the months in the period). The annual and interview trends might be expected to show the most divergence in these two periods. A further note on the 2008/9 annual rate is that March 2009 is not covered in the data (the last available month of interviews was in that month, providing a reference period up to the February). A comparable measure was produced by taking the mean of the previous 11 months and using this as the rate for March.

The trends are plotted against one another in line with the central point in their reference period. The central reference point of 2005/6, for instance, would be March 2005 – around six months behind the centre of the annual 2005/6 measure. In this way, we eliminate time lags to the extent that we can in the presentation of these trends.

### Comparing the trends

Figure 4.5 displays the trends for 'all vandalism'. Analysing the trend comparison in some detail for this crime type will provide a basis for interpreting trends in the other types of crime presented in the charts.

Overall, there is a very close relationship between the measures. Rates are both at a very similar level and follow a similar trend, particularly in the middle of the period.

Looking more closely, there are some points of divergence to be explained. In the earliest period, the measure of annual crimes provides an estimate for 2001/2 that is well below that for the interviews in that period. As noted above, the 2001/2 estimates will be particularly susceptible to sampling error due to the low numbers of interviews available for the earliest quarter in particular (the sample size for the April 2000 monthly estimate was 1,428 compared with 34,733 for the estimate one year later). The annual measure gives equal weight to the monthly estimates. Looking at the estimates for those earliest months, we do find that the three lowest rates are found in the first four months. Although these respondents will be included in the interviews measure, their influence on the overall estimate will be much lower.

A second point of noticeable divergence is around the peak of the trend for the interviews measure at 2006/7 (this represented a statistically significant increase from the previous year). It was noted in the previous section that the interviews measure would in general produce a smoother trend, but here the interviews measure peaks above the line for the annual measure. The explanation for this lies in the different mid points in the measures presented. Although not presented in this report, it was possible to look at the monthly rates used for the annual measure to provide some insight into crime levels at particular points. It was observed that in the eight months that form the centre of the reference period for the 2006/7 interviews measure, the monthly rates were higher on average than for the months either side. The result is a peak in the interviews measure that is not replicated in the annual measure simply because the two do not share the same mid-points in their reference periods. Instead, a relatively high rate is maintained across two years with the annual measure whilst there is a fall in the interviews estimate. An annual estimate that related to October 2005 to September 2006 would have provided a similar peak to that seen in the interviews measure.

This analysis of the vandalism trends provides explanation for the similarities and divergences found in the trends for other crime types. In general, the trends are very similar between the measures. Burglary and theft from vehicles are very close indeed, as are the all violence and all personal crime trends (except at the beginning and end of the period where sample sizes are relatively small for the annual measure). Quite a large component of the 'all household crime' rate relates to vandalism, and some of the same features of the trends for that crime type discussed above can be observed.

The theft of vehicles and theft from the person trends (both relatively low incidence crimes) reflect the smoother interviews trend predicted by the analysis with simulated data above. This is also the case with the trends for violence with and without injury, although here there is also evidence of the effect of there being different mid-points in the reference periods.

Figure 4.5 Vandalism – trends for annual crime and years of interviews (rates per 10,000)






Figure 4.6 Burglary – trends for annual crime and years of interviews (rates per 10,000)







**Figure 4.7 Burglary with entry – trends for annual crime and years of interviews (rates per 10,000)**



**Figure 4.8 All vehicle-related theft – trends for annual crime and years of interviews (rates per 10,000)**



**Figure 4.9 Theft from vehicles – trends for annual crime and years of interviews (rates per 10,000)**



**Figure 4.10 Theft of vehicles – trends for annual crime and years of interviews (rates per 10,000)**




Figure 4.11 All household crime – trends for annual crime and years of interviews (rates per 10,000)




Figure 4.12 All personal crime – trends for annual crime and years of interviews (rates per 10,000)




Figure 4.13 Theft from the person – trends for annual crime and years of interviews (rates per 10,000)




Figure 4.14 All violence – trends for annual crime and years of interviews (rates per 10,000)




Figure 4.15 Violence with injury – trends for annual crime and years of interviews (rates per 10,000)




Figure 4.16 Violence without injury – trends for annual crime and years of interviews (rates per 10,000)





Figure 4.17 Domestic violence – trends for annual crime and years of interviews (rates per 10,000)



## 4.5 Conclusions

### Closely comparable trends

The closeness of the trends observed between the two types of measure means there is little to suggest that the Home Office should be concerned that the interview years approach it has adopted is under- or over-stating movements in crime rates. There will be a small degree of smoothing in comparison with the annual measure, although this is often overridden by the different lengths of reference periods and different mid-points.

### Reporting issues

We understand that it is not intended at this stage to introduce an annual measure of crimes over the current arrangement in reporting, but rather to use the comparison to assist with the interpretation of the 'year of interviews' measure.

A reason to employ a measure of *crimes* in a year rather than *interviews* in a year is that the former allows for a more intuitive understanding of the measure (even if the actual calculation is no more straightforward). With the latter, describing the time period in which crimes fall is not a simple matter. To date, the description of this measure has referred to "crime according to 2007/8 interviews". An accurate (if less snappy) description of the same estimate might refer to "crime in the two years centred on March 2007", although this again does not properly describe the higher weight given to the central months in the distribution. By comparison, it may be more attractive to have a measure that could be understood, for instance, as "crime in 2007/8".

A point to note for the future comparison of annual crime and year of interview measures is the additional six months of data required for the former measure. For an estimate based on *interviews* in 2007/8, an estimate for *crime* in the year October 2006 to September 2007 is the closest comparison (in terms of centre of reference period), but this period is covered in interviews right up to September 2008 – six months after the interviews measure. For this reason, direct comparisons could not be made for the latest estimates without incurring delays in publication (although estimates could be produced for the latest period based on the data available by that point and could be revised in later publications).

### **Improving the annual crime measure: weighting and significance testing**

For the annual rates to be published alongside the year of interview data on a regular basis it may be desirable to look again at the options for recalculating the calibrated weights. This would not be an easy task to work through in the first instance and would be time-consuming on a regular basis, so may not be considered worthwhile. Similarly, an approach to significance testing could be developed that allowed the reader to be confident about the status of changes in rates over time. Again, this would be a challenging task and may be disproportionate depending on the use to be made of the alternative measure.

## Appendix A Chapter 2: additional tables

Table A.1 Sample sizes (actual) for inner city areas and metropolitan PFAs								
	Total core sample	Inner city sample	Other core sample	% sample inner city	London Met sample	% sample London	Greater Manchester sample	% sample Greater Manc
2008-09	44,002	3,629	40,373	8%	3,693	8%	1,341	3%
2007-08	46,983	3,699	43,284	8%	3,632	8%	1,570	3%
2006-07	47,203	3,675	43,528	8%	3,526	7%	1,543	3%
2005-06	47,796	3,736	44,060	8%	3,367	7%	1,535	3%
2004-05	45,120	3,641	41,479	8%	3,347	7%	1,375	3%
2003-04	37,931	3,395	34,536	9%	3,397	9%	1,543	4%
2002-03	36,479	3,219	33,260	9%	3,376	9%	1,414	4%
2000	19,410	2,363	17,047	12%	2,186	11%	880	5%
1998	14,947	3,542	11,405	24%	2,385	16%	750	5%
1996	16,348	3,870	12,478	24%	2,559	16%	837	5%



**Table A.2 Crime measures used in the analysis**

PERSONAL	HOUSEHOLD
Sexual offences	Vandalism
Against women	Motor vehicle vandalism
Against men	Other vandalism
Common assault	Burglary
Wounding	Attempts
Robbery	Attempts and no loss
Theft from person	With entry
Theft from person & robbery	With loss
Mugging	Theft from a dwelling
Other thefts of personal property	Theft from a motor vehicle
Comparable personal	Theft of a motor vehicle
All personal (including sex)	Attempted theft of & from vehicle
All personal (not including sex)	All vehicle thefts
All assault (common + wounding)	All vehicle crime
Threats	Bicycle theft
	Other household thefts
All BCS violence (includes snatch theft)	Comparable household
Domestic violence	All household offences
Mugging	
Stranger	
Acquaintance	
Comparable violence	
Comparable household	
All household offences	
Comparable personal (not sex)	
All personal (not including sex)	
Acquisitive crimes	
Vandalism	
Robbery and Wounding	
Stealth theft from person	
Snatch theft from person	
Mugging + stranger	
Domestic + acquaintance	
Household acquisitive	
Personal acquisitive	
Assault with minor injury	
Assault without injury	
BCS all violence (no snatch theft)	

**Table A.3** Design factors attributable to disproportionate selection probabilities, clustering and weighting for national estimates

BCS year	Actual sample size	Effective sample size	Sample efficiency	Design factors of the core sample
2008-09	44,002 <sup>1</sup>	26,037	59%	1.30
2007-08	46,983	27,452	58%	1.31
2006-07	47,203	27,647	59%	1.31
2005-06	47,796	25,837	54%	1.36
2004-05	45,120	25,457	56%	1.33
2003-04	37,931	21,034	55%	1.34
2002-03	36,479	21,490	59%	1.30
2000	19,410	11,935	61%	1.28
1998	14,947	9,017	60%	1.29
1996	16,346	10,508	64%	1.25

<sup>1</sup> Each survey year contains a small number of cases rolled forward from the previous one. There were 2,284 cases in the 2008-9 survey data that had been rolled forward from 2007-08. These were excluded for the purposes of evaluating the effects of the sample design as they would have been selected under the previous design.

**Table A.4** Design factors attributable to disproportionate selection probabilities, by PFA

Please refer to Excel workbook: Table A4\_4xA4pages.xls

## Appendix B Chapter 2: the impact of increasing the sample size on PFA level estimates

The increase in sample size will have improved the precision of estimates and allowed smaller differences to be detected between different years. Power calculations indicate what size differences can be detected in key estimates between different PFAs and within the same PFAs over time. Power calculations show how large the difference between two estimates needs to be before we are able to state that the difference is statistically significant, i.e. it is genuine and not there by chance.

The effective sample size has been used to assess the increase in power, as this more accurately reflects the level of precision (shown in Table B.1). BCS 2004-5 and onwards had a minimum sample size of 1000 per PFA and an average effective sample size of 700 per PFA (excluding Metropolitan and Greater Manchester, which contain considerably more interviews than other PFAs). From 2001-2 until 2003-4 a minimum of 700 was set per PFA, which results in an average effective sample size of 550. In 2000 there was a minimum of 300 per PFA and an average effective sample size of 325 (since the majority of PFAs achieved far more than the minimum sample size).

As well as the sample size, power is affected by the size and variance of the estimate. Crimes with a lower prevalence will have lower power, as will estimates with a large variance. The variance for the crime estimates is generally fairly high, this is due to the nature of crime estimates; the majority of individuals have a low prevalence but there is a core group who experience a large amount of crime.

As a rule of thumb a test with power of 0.8 or more is usually desired. A power of 0.8 means the probability of not detecting a real significant difference is 20%. This level of power was achieved for crimes with a higher prevalence and for larger differences between two estimates. Table B.2 shows the power for a set of example estimates which range from 0.025 (a prevalence rate of 250 per 10,000) to 0.125 (1250 per 10,000). It shows the power of a statistical test between the example estimate (Estimate A) and a second estimate (Estimate B) that is 7%, 10% or 3% higher or lower. Power is given for three samples with effective sample sizes of 325, 550 and 700.

A statistical test of the difference between an estimate of 0.1 and an estimate of 0.1130 (13% higher than the initial estimate) both from a PFA with an effective sample size of 550 would have a power of 0.82, meaning we would be confident that the statistical test is picking up a genuine difference. A statistical test of the difference between an estimate of 0.05 and an estimate of 0.0565 (13% higher) both from a PFA with an effective sample size of 325 would have a power of 0.43, meaning there is a 57% probability that we are not detecting a real significant difference.

**Table B.1 Effective sample sizes by PFA**

<b>Year</b>	<b>2008-09</b>	<b>2004-08</b>	<b>2001-04</b>	<b>2000</b>	<b>1998</b>	<b>1996</b>
<b>Minimum number of interviews per PFA</b>	<b>1000</b>	<b>1000</b>	<b>700</b>	<b>300<sup>1</sup></b>	<b>N/A<sup>2</sup></b>	<b>N/A<sup>2</sup></b>
Metropolitan	2667	2368	2119	1507	1566	1655
Greater Manchester	1027	1070	988	528	580	680
Merseyside	728	707	570	281	293	407
South Yorks	780	779	570	440	312	403
Northumbria	850	617	609	410	217	486
West Midlands	845	1066	901	666	616	712
West Yorks	890	904	725	634	524	400
Avon & Somerset	700	727	591	289	296	314
Bedfordshire	782	668	443	249	163	147
Thames valley	464	680	693	672	528	590
Cambridgeshire	625	655	495	262	264	236
Cheshire	587	715	530	264	237	265
Cleveland	686	568	496	226	97	80
Devon & Cornwall	610	672	603	567	331	418
Cumbria	739	665	509	303	117	160
Derbyshire	703	731	538	312	201	199
Dorset	590	651	436	201	25	68
Durham	865	651	550	318	253	135
Sussex	724	659	433	432	333	489
Essex	759	728	526	320	270	282
Gloucestershire	800	632	464	239	165	163
Hampshire	722	703	592	432	347	398
West Mercia	712	640	486	425	272	147
Hertfordshire	939	639	560	239	190	163
Humberside	478	648	352	194	232	257
Kent	780	752	526	408	426	342
Lancashire	496	686	646	368	277	262
Leicestershire	799	634	504	232	250	256
Lincolnshire	583	681	662	206	161	227
Norfolk	663	627	560	193	182	224
Northamptonshire	731	636	449	204	74	159
North Yorks	733	591	434	263	105	147
Nottinghamshire	659	665	442	281	202	310
Staffordshire	572	708	569	276	299	335
Suffolk	473	763	637	271	121	165
Surrey	749	699	578	258	58	238
Warwickshire	807	685	630	267	129	181
Wiltshire	796	739	537	252	188	188
North Wales	514	705	525	172	200	158
Dyfed Powys	584	759	507	229	88	69
Gwent	828	704	568	345	64	205
South Wales	923	707	505	347	229	345
Average (excl Met and Greater Manchester)	707	696	549	324	233	268

<sup>1</sup> Although BCS 2000 set a min of 300, the achieved sample size in most PFAs was higher than this, giving an average actual sample size of 400 per PFA.

<sup>2</sup> PFA was not used as a stratification variable and there was no minimum sample size set.

**Table B.2 Power for differences between two estimates with the same sample size**

Estimate A	Difference between estimates	Value of Estimate B	Power of test from a sample with an effective sample size of...		
			700	550	325
0.025	13%	0.0218 or 0.0283	0.41	0.40	0.25
0.050		0.0435 or 0.0565	0.70	0.57	0.43
0.100		0.0870 or 0.1130	0.91	0.82	0.65
0.125		0.1088 or 0.1413	0.96	0.95	0.74
0.025	10%	0.0225 or 0.0275	0.26	0.26	0.17
0.050		0.0450 or 0.0550	0.48	0.37	0.28
0.100		0.0900 or 0.1100	0.72	0.60	0.43
0.125		0.1125 or 0.1375	0.82	0.79	0.52
0.025	7%	0.0233 or 0.0268	0.15	0.15	0.11
0.050		0.0465 or 0.0535	0.27	0.21	0.16
0.100		0.0930 or 0.1070	0.43	0.34	0.24
0.125		0.1163 or 0.1338	0.52	0.49	0.29

Many of the crime prevalence rates are low, which decreases the power of the tests, and increases the possibility of results being falsely negative. This means it is possible that a difference between the crime estimates of two PFAs will not be significant in statistical tests when it really is. However, it is unlikely that statistical tests will give false positives, meaning it is unlikely that a false significant result will be returned. This makes the statistical tests conservative in nature.

## Appendix C Chapter 2: additional information on analysis of the impact of changes to the degree of clustering

The values for rho have been estimated directly, rather than having been taken from the deffs published in technical reports. This is because we wished to isolate the clustering effects from other design effects for this specific analysis.

The intra class correlation coefficient (rho) has been used to examine the effects of clustering. Rho measures how homogenous the respondents within a cluster are and is estimated as the ratio of between cluster variance and total variance:

$$\rho = \frac{\sigma_b^2}{\sigma_b^2 + \sigma_w^2}$$

Where  $\sigma_w^2$  is the within cluster variance and  $\sigma_b^2$  is the between cluster variance.

The value of rho ranges between zero and one. A value of one implies that all the respondents within the cluster are identical, whereas a value of zero means there are no similarities and all the respondents are unique. Respondents are expected to be more alike when clusters are small, subsequently we would expect highly clustered samples to have larger values of rho. We would expect a completely un-clustered sample to have a rho of zero.

Rho was estimated for each of the key crime estimates for all years of BCS from 1996 onwards. The estimates were carried out in Stata version 10 by fitting a 2-level random effects model. The estimates of rho were used to calculate design effects. Design effects were generated for each of the key crime estimates outlined in Table A.2. The design factors have been estimated using the following formula:

$$deff \approx 1 + (m - 1)\rho$$

Where  $\rho$  is the intra class correlation coefficient and  $m$  is the average number of interviews per cluster, shown in Table C.1.

Table C.1	Average number of achieved interviews per PSU				
	1996	1998	2000	2001-8	2008-9 <sup>1</sup>
Overall	20	23	21	21	-
Inner city	20	23	19	16	-
Other areas	20	23	22	21	18

<sup>1</sup> Inner city areas not clustered

### Prevalence

Earlier survey years had more incidents of crime. Table C.2 shows the unweighted prevalence of total BCS crime for each of the survey years. It shows inner city areas have a higher prevalence of crime.

Table C.2 Unweighted crime prevalence by survey year					
	1996	1998	2000	2001-8	2008-9
Unweighted prevalence for total BCS crime	38%	33%	29%	23%	21%
<i>Other areas</i>	36%	31%	28%	22%	20%
<i>Inner city</i>	44%	39%	33%	28%	25%

### Data from 2001-2008

To maximise the sample sizes available for analysis, a number of years of data have been grouped together. The 1996, 1998 and 2000 surveys used quarter postcode sectors as PSUs and have been grouped together. The 2001-2 to 2007-8 surveys used whole sectors and have also been grouped. The 2008-9 survey used a partially un-clustered design and has been kept separate.

The continuous fieldwork used on the BCS since 2001 means it is possible for the addresses in a single cluster to end up in two different years of data. For example, a PSU issued near the end of the fieldwork period may have had half the cases worked in time to be included in the current year of data, but data from the remaining cases would be included in the following year. This feature made it necessary for all data from 2001-08 to be merged and analysed together. It was felt that crime rates had been steady enough over these years to allow this.

## Appendix D Chapter 3: additional tables

Table D1 – see Spreadsheet File: Table\_D1\_v2.xls

Table D2 – see Spreadsheet File: Table\_D2.xls

Table D3 – see Spreadsheet File: Table\_D3\_v2.xls



# Appendix E Chapter 4: additional tables

**Table E.1 Trends for annual crime and years of interviews (rates per 10,000)**

*Base: All individuals / households*

*BCS*

Offence categories (annual and interview measures)	Year in which incidents occurred / year of interviews								
	2000	2001/2	2002/3	2003/4	2004/5	2005/6	2006/7	2007/8	2008/9
	Rates per 10,000 population / households								
Vandalism (annual)	1036	1157	1136	1091	1170	1231	1232	1182	1165
Vandalism (year of interviews)	-	1185	1145	1104	1125	1182	1281	1141	1161
Burglary (annual)	456	436	438	382	310	315	312	321	304
Burglary (year of interviews)	-	441	439	422	331	317	311	312	312
Burglary with entry (annual)	251	255	249	223	195	186	184	193	180
Burglary with entry (year of interviews)	-	251	253	239	205	190	182	186	190
Vehicle-related theft (annual)	1117	1096	1027	880	761	738	706	617	631
Vehicle-related theft (year of interviews)	-	1135	1068	950	827	749	723	639	635
Theft from vehicles (annual)	693	645	642	558	487	493	467	421	447
Theft from vehicles (year of interviews)	-	681	643	599	531	485	483	424	444
Theft of vehicles (annual)	149	145	110	97	91	70	74	68	54
Theft of vehicles (year of interviews)	-	144	126	108	94	80	75	69	63
All household crime (annual)	3511	3482	3351	3061	2942	2986	2938	2830	2760
All household crime (year of interviews)	-	3586	3428	3217	2978	2939	3038	2732	2831
Theft from the person (annual)	116	159	161	141	123	139	128	163	170
Theft from the person (year of interviews)	-	144	164	147	137	134	132	132	164
All violence (annual)	528	674	629	566	536	565	538	480	555
All violence (year of interviews)	-	652	644	612	544	545	569	502	478
Violence with injury (annual)	310	348	349	301	299	282	273	245	281
Violence with injury (year of interviews)	-	358	342	321	305	285	292	242	252
Violence without injury (annual)	218	325	280	265	237	283	265	235	274
Violence without injury (year of interviews)	-	294	302	291	239	260	276	259	225
Domestic violence (annual)	96	140	110	98	90	78	87	69	76
Domestic violence (year of interviews)	-	150	120	105	94	83	94	78	66
All personal crime (annual)	988	1157	1128	979	927	975	914	877	986
		1132	1127	1071	951	956	963	859	889

Notes:

1. 2008/9 includes only months April 2008 to February 2009 for annual crime but the estimate in the table multiplies the estimate for the 11 months up to 12 months using the average of the 11 months.
2. Estimates for the annual measure are about six months later than for the year of interviews measure (e.g. central point for annual measure for 2007/8 is Sept/Oct 2007 compared with March 2007 for the year of interviews measure).

Table E.2 Unweighted sample sizes for the estimates of crime for each month									
	2000/1	2001/2	2002/3	2003/4	2004/5	2005/6	2006/7	2007/8	2008/9
Apr	1,428	34,733	36,600	37,436	45,681	47,945	47,344	47,567	41,888
May	2,906	36,922	35,570	38,459	45,154	48,972	47,253	47,335	38,051
Jun	4,738	37,867	36,068	39,211	44,290	49,905	47,498	46,654	34,391
Jul	7,398	38,490	36,497	39,753	44,340	50,004	47,109	47,336	29,695
Aug	10,670	38,513	36,629	40,383	45,358	48,975	47,376	46,640	26,078
Sep	13,693	38,460	36,881	40,953	46,011	48,263	47,905	45,839	22,617
Oct	16,851	38,291	36,767	41,647	46,761	47,990	48,342	45,455	18,518
Nov	19,749	38,042	37,021	42,500	46,951	48,526	47,894	44,977	14,962
Dec	22,589	37,786	37,327	42,672	47,239	48,170	48,028	44,580	12,231
Jan	26,944	36,742	37,203	44,071	47,325	47,733	47,760	45,377	7,467
Feb	30,279	35,875	37,518	44,920	47,769	47,308	46,740	46,511	3,702
Mar	32,824	36,479	37,931	45,120	47,796	47,203	46,983	46,286	

# Appendix F Chapter 4: issues for an alternative measure of prevalence rates

## The problem of developing an alternative prevalence measure

Chapter 4 of this report focused on the development of an alternative measure of incidence rates that was focused on incidents that occurred in a particular period. The calculation of prevalence rates for a particular year presents a different and more challenging problem to that of incidence rates.

Prevalence is specifically an *annual* measure *for an individual*, and its property as a binary measure adds to the difficulty.

The problem with the continuous survey design is that for most sample members only part of their reference period will be within the required year. Only interviews in one month will provide a measure of the entire period for individuals, and any prevalence measure based just on a single month of interviews would not provide the required precision of estimates.

Although we could construct a *monthly* prevalence rate for individuals in different monthly samples, it is not then possible to calculate an annual rate by adding these together as we have done for incidence rates. The annual prevalence needs to be measured *at the individual level for that period*. For instance, a prevalence of an offence of three per cent in each of 12 months does not necessarily result in an annual prevalence of 36 per cent – repeat victimisation means that it is likely that the actual annual prevalence would be considerably lower.

This is demonstrated in Table F.1. In that table, data is provided for ten cases over two months in the left hand columns (columns 'a' and 'b'). The prevalence across the two months is calculated (seven out of ten) and that of the months individually (five and four out of ten). Clearly, there is less chance of experiencing crime in one month than there is in two months, and for this reason taking either month individually as an estimate for the two month period underestimates the actual prevalence (five or four rather than the actual seven). Further, adding the prevalence together for the two months gives a figure of nine - an overestimate due to repeat victimisation in cases F and G.

It might be considered that some kind of weighting approach might be used to account for individuals whose relevant reference period is less than the full 12 months. However, we again come up against the problem of repeat victimisation. In Table F.1, a weighted approach is presented in the last two columns where the prevalence in one month is multiplied by two to reflect the two months in the required prevalence period. The result is an overestimate of prevalence even when averaging between the two weighted estimates because in some cases the individual has experienced more than one crime in the period. A further problem with the approach described is that it loses its intuitive appeal to have a prevalence of 'two' for a measure that can only be victim or non-victim.

**Table F.1 Example of prevalence calculations over a two month period**

	Victim (1) or not victim (0) in period		Prevalence			Weighted prevalence (prevalence for the month multiplied by the number of months required)	
	Month 1 a	Month 2 b	Months 1 and 2	Month 1	Month 2	Month 1 a*2	Month 2 b*2
A	1	0	1	1	0	2	0
B	1	0	1	1	0	2	0
C	1	0	1	1	0	2	0
D	0	1	1	0	1	0	2
E	0	1	1	0	1	0	2
F	1	1	1	1	1	2	2
G	1	1	1	1	1	2	2
H	0	0	0	0	0	0	0
I	0	0	0	0	0	0	0
J	0	0	0	0	0	0	0
<b>Total</b>			<b>7</b>	<b>5</b>	<b>4</b>	<b>10</b>	<b>8</b>

## Possible solutions: time series modelling and matching sample members in different offence periods

At this stage, perhaps the most promising approach to a solution would be to produce a series of estimates by month of survey (with the reference period corresponding to that asked in the survey) and then use time series methods to combine these monthly estimates to produce an estimate for an annual period. However, this would be a complex undertaking and is beyond the scope of the current study. Furthermore, an approach to producing annual estimates that relies on modelling is will result in a lack of certainty about whether any divergence between trends in the interview year and annual prevalence measures was simply the result of imperfections in the model. In any case, we would not have produced an intuitive measure, which was the major aim of the task.

A further possible solution may lie in making use of the property of the distribution of interviews that means that, broadly speaking, for each individual whose reference period only covers one of the 12 months of a period of interest, there is someone else whose reference period covers 11 of them. We cannot simply randomly join together the victimisation history of a respondent appearing on one side of the offence period with another appearing on the other as victimisation and repeat victimisation are not random events. A solution may lie in simulating the prevalence for an individual by modelling victimisation and then matching nearest neighbours who appear in different points in the offence year. Prevalence would then be calculated for the 'simulated' individual. A concern for this approach may be that it would not adequately capture the degree of persistence of victimisation for a given individual, and it suffers from the same problem noted above of not producing an intuitive measure.

## The development of an annual prevalence measure

If the aim of creating the annual measures of crime victimisation was to produce intuitively appealing estimates that could be compared to the current trends to check the story that they present, it would not seem worthwhile to develop a complex model-based set of prevalence estimates. Even if measures could be produced (for instance using a time series approach) that we could have

confidence in, it would seem unlikely that they would be sufficiently accurate to provide a check on the current prevalence trends, and would certainly not be intuitive.