**Statistics
New Zealand**
TATAURANGA AOTEAROA

Creating regional crime statistics from administrative data

Statistics New Zealand Working Paper No 13–02

Rebecca McGirr, John Bryant, and Tim Hawkes

New Zealand Government

Disclaimer

The Statistics New Zealand Working Paper series is a collection of occasional papers on a variety of statistical topics written by researchers working for Statistics New Zealand. Papers produced in this series represent the views of the authors, and do not imply commitment by Statistics NZ to adopt any findings, methodologies, or recommendations. Any data analysis was carried out under the security and confidentiality provisions of the Statistics Act 1975.

Liability statement

While all care and diligence has been used in processing, analysing, and extracting data and information in this publication, Statistics New Zealand gives no warranty it is error free and will not be liable for any loss or damage suffered by the use directly, or indirectly, of the information in this publication.

Contents

List of tables and figures

Tables by section

Figures by section

Abstract

Objectives

The only publicly-available information on the geographical distribution of crime in New Zealand is offence statistics for police administrative units. We investigate whether existing data can be used to construct geographical crime statistics that correspond to regional councils, territorial authorities, and urban areas.

Methods

We build experimental output geographies from police stations, the smallest administrative unit for which there are long time series of offence statistics. We develop three rules for assigning police stations to the new geographies: one based on population, one based on land area, and one based on both. We assess the performance of these rules by calculating the proportion of national land area and population that is misclassified, and the number of target units that do not receive at least one police station. We also look at whether regional statistics on serious assaults are sensitive to the choice of allocation rule.

Findings

The new output geographies approximate the target geographies well. For instance, our preferred rule assigns 96 percent of the national population to the correct territorial authority. Moreover, a case study of serious assaults suggests that most regional crime statistics are not sensitive to the choice of rule.

Conclusion

The new output geographies perform sufficiently well that they could, if required, be used to produce regional crime statistics.

Key words

crime, offence statistics, geography, administrative data

Acknowledgements

1   Introduction

Figures disprove Christchurch's crime reputation

Christchurch's reputation as the country's violent crime capital has been contradicted by the latest police figures. … The figures do not provide direct comparisons of city crime rates, but total violent crimes recorded in Christchurch central last year (1,465) were the lowest of the main centres.

The Press, 3 April 2010

Many New Zealanders are interested in how crime rates in their neighbourhood compare with rates in other parts of the country. At present, the main public source of information on geographical variation in crime is offence statistics for police areas and districts, available on the Statistics New Zealand and New Zealand Police websites. These are the numbers most commonly reported by the media, including the Press article cited above. Police areas and districts are suitable for some sorts of comparisons, such as examining police workloads or performance. But to answer questions about which areas have the most crime, the geographical units that are used should align with standard definitions for these areas.

Figure 1

Christchurch population density and police areas

Source: Calculated from 2006 Census data on usually resident population by meshblock.

The size and shape of police districts and police areas reflect the particular administrative needs of the police. They do not generally align with standard geographical definitions. Figure 1 gives one example. The map shows police areas overlaid on population density. The Christchurch Central police area is the unit most commonly used for describing crime in Christchurch. People familiar with Christchurch will recognise that the Christchurch Central police area takes in the central business district, Hagley Park, some inner suburbs, and wetlands. Most of the city is omitted. Even before the earthquakes of 2010 and 2011, crime statistics for the Christchurch Central police area are unlikely to have been a reliable guide to crime across the city as a whole.

The rest of the country also has notable misalignments between police areas and districts and the regional council and territorial authority boundaries. Figure 2 demonstrates the misalignment between police and territorial authority classifications. Territorial authorities are mapped as coloured shapes with police district and area boundaries overlaid as black and yellow borders, respectively.

Figure 2

Territorial authorities with police district and area boundaries

Source: Statistics New Zealand

Compounding these problems is the absence of standard, well-documented output geographies for crime statistics. This means that users of these statistics often do not know which geographical units are being referred to.

In New Zealand, the police station is the smallest geographical unit we have comprehensive time series of offence statistics for. Detailed, consistent electronic data at the police station level exist back to 1994, permitting trends to be analysed over almost two decades. By grouping police stations into larger units that approximate standard geographies such as territorial authorities and regional councils, it is possible to estimate offence rates for these geographies. One study that takes this approach is the *Quality of life in twelve of New Zealand's cities 2007* report (Quality of Life Project 2007: 91–98), which presents offence rates for burglaries, violence, sexual offences, car offences, and drug and antisocial offences at the territorial authority level. Another is the 'Regional indicators' section of the Ministry of Social Development's Social report, which presents rates for all offences at the regional council level.

Aggregating police station data is a promising approach to creating time series for regional crime statistics. However, any attempt to produce general-purpose statistics from administrative data inevitably requires trade-offs between accuracy, timeliness, and transparency (Rees 1986; Freedman et al 2008; Gregory et al 2010). It is important to assess these tradeoffs, to help inform methodological choices, and to provide users with information about data quality.

In this paper we investigate three simple rules for constructing regional crime statistics out of police station data. The rules assign police stations to target geographies based on population, land area, and a combination of the two. The target geographies that we use are regional councils, territorial authorities, and main urban areas. We evaluate our new geographies by calculating the extent to which population and land area are misclassified, and by examining whether different rules for allocating police stations have a material effect on regional crime comparisons.

We find that the allocation rule that uses both population and land area performs best. The difference in performance is generally small, however, with all three rules performing well, except for smaller urban areas. We suggest that it would be feasible to construct crime statistics for regional councils, territorial authorities, and large urban areas, but not for small urban areas.

2 Geographical classifications

2.1 Geographical classifications used by the police

Summary information about the police districts, areas, and stations is provided in table 1. The boundaries of the police districts are shown in figure 3. The geographical classification is hierarchical, so that any police station falls entirely within a single police area and any police area falls within a single police district. As can be seen from figure 3, there is enormous variation in the size of the units. Data from the 2006 Census (not shown) indicate that population sizes are equally variable, with the smallest police station covering 140 people, and the largest covering 120,000.

Table 1

Area and population statistics for police classifications

Classification	Number	Mean area (thousands of km^2)	Mean population (thousands of people) in 2006
Police geography			
District	12	34.6	360
Area	43	9.7	100
Station	287	1.4	15
General purpose geographies			
Regional council	16	25.8	270
Territorial authority	73	3.8	59
Main urban area (zones separated)	25	0.2	125
Main urban area (zones combined)	16	0.3	195
Source: 2006 Census of Population and Dwellings			
Note: The results for numbers, areas, and populations do not include residual categories such as 'Area outside regional council' and 'Area outside territorial authority'.			

The New Zealand Police National Recording Standard (2008: 25) requires that two types of police station are recorded for each offence that is reported. The first is the 'scene station', which is the station where the offence occurred. The second is the 'reporting station', which is the station where the occurrence was reported. In this paper, station always means scene station.

Figure 3

Police areas and police districts

Source: Statistics New Zealand

Note: Each police district is represented by a different colour.

2.2 General purpose geographical classifications

Regional councils and territorial authorities are important administrative units. Most official statistics below the national level use classifications based on their boundaries. On 1 November 2010 the seven territorial authorities making up the Auckland region were amalgamated into a single Auckland Council. In this paper, however, we use the pre-amalgamation territorial authorities. This provides evidence on the feasibility of providing statistics for areas within Auckland, and maintains historical continuity.

Main urban areas, in contrast, are a statistical rather than an administrative unit. They are predominantly based on data from the 1991 Census. The urban classification is designed to "identify concentrated urban or semi-urban settlements without the distortions of administrative boundaries" (New Zealand Standard Areas Classification 1992, p28). The Auckland, Wellington, Hamilton, and Napier-Hastings main urban areas are subdivided into zones. For some purposes, such as comparing between main centres, it is useful to have data at the level of the main urban area. For other purposes, such understanding trends in particular cities, it is useful to have data at the level of the zone. We experiment with two different types of urban area classification: one that distinguishes between zones and one that does not.

3 Data and methods

3.1 Data

All mapping and geographical analysis in the paper is based on the 2010 New Zealand Transverse Mercator digital boundaries. ESRI shapefiles and Mapinfo TAB files with these boundaries can be downloaded from the Statistics NZ website.

The unit for much of our analysis is the meshblock, the smallest element in New Zealand's geographical hierarchy. We assembled a concordance file showing, for almost every meshblock in the country, the regional council, territorial authority, urban area, and police station to which each meshblock belongs. Police stations are based on boundaries as of 28 January 2009. All other geographical units were based on 2010 boundaries. The small number of meshblocks omitted from our concordance file all lie outside the official boundaries of New Zealand police stations.

We added estimates of land area and the geographical distribution of the population to the concordance file. The digital boundaries file contains estimates of the size of each meshblock. To obtain the land area, we excluded 'watery' meshblocks (ones where the 'iwtext' field equalled 'inland water', 'inlet', 'oceanic', or 'other'). We obtained meshblock-level estimates of the usual resident population from the 2006 Census. Many meshblocks had been split into two or more smaller meshblocks between 2006 and 2010. We allocated the population of the split meshblocks equally across the newly created units. More sophisticated allocation rules could be devised, but it is unlikely that they would have had a material effect on our results.

Statistics NZ's Population Statistics Unit gave us a tabulation with annual figures for the estimated resident population by police station between 2001 and 2008. The estimates refer to 30 June of each year, and are based on the 28 January 2009 police station boundaries. Population estimates for earlier years are not readily available, and are not essential for testing the allocation rules, so we restrict the analysis to 2001 and later.

The New Zealand Police provided customised tabulations of recorded violent offences, by offence class, police station, and calendar year, for 1994–2008, though we only use data for 2001 to 2008.

3.2 Rules for assigning police stations to target geographies

We have data on offences disaggregated to the level of the police station. We have 'target geographies'—regional councils, territorial authorities, and main urban areas—whose offence rates we wish to approximate. We need rules for allocating individual police stations, and hence offences, to units within the target geographies.

Figure 4 provides an example of the ambiguous cases that an allocation rule must resolve. The map shows police stations in and around Christchurch, plus the Christchurch city and Selwyn district territorial authorities. The question is how police station data should be used to approximate offence rates in the Christchurch city territorial authority. Some stations, such as New Brighton and Sydenham, fall entirely within Christchurch city. Any sensible allocation rule will assign offences committed in these stations to Christchurch city. Other stations, such as Hornby, lie across two territorial authorities. In contrast to New Brighton and Sydenham, it not clear how to allocate offences committed in the Hornby police station area to the Christchurch city and Selwyn district territorial authorities.

Figure 4

Police stations and territorial authorities in Christchurch

Source: Statistics New Zealand

Note: The green patches denote territorial authorities, and the lines denote police stations. The Hornby police station, for instance, is located mainly in the Christchurch city territorial authority, but lies partly within the Selwyn territorial authority.

One way of dealing with a station that falls across multiple units of the target geography is to associate the station with all of these units and divide offences for that station among the units using some set of weights. For instance, Hornby police station could be assigned to both Christchurch city and Selwyn district, with, perhaps, 90 percent of offences being allocated to Christchurch central and 10 percent to Selwyn district. An alternative is to allocate each overlapping station to a single target unit. For instance, Hornby police station could be assigned entirely to Christchurch city, so that all offences recorded to Hornby police station were counted as occurring within the Christchurch city territorial authority.

Allocating overlapping police stations to multiple units of the target geography might reflect actual patterns more accurately than allocating them to a single unit, if the weighting system reflected the true distribution of crime within each police station. However, allocating stations to multiple units is more complex, in that it requires more information and more calculations. This makes it harder to replicate the resulting statistics or to explain the procedures to a non-technical audience. Allocating stations to multiple units also implies reporting fractional numbers of offences. For instance, it might lead to a report that 0.9 murders were committed in Christchurch city and 0.1 in Selwyn district. These counter-intuitive results can undermine the credibility of the statistics.

Simplicity, transparency, and credibility are important virtues for all official statistics, but particularly for statistics that attract intense interest from non-technical audiences. We have therefore restricted our investigation to allocation rules that assign overlapping stations to single units of the target geography.

Having decided to assign each police station to a single unit, we still need to decide on criteria for making the assignments. One obvious possibility is land area. Looking at figure 4, for instance, it is natural to conclude that Hornby should be allocated to Christchurch city because most of its land area is within Christchurch city. The decision would be less clear cut, however, if most of the population of Hornby happened to live in the corner of the police station that lies within Selwyn district. Given that offences are committed by people, it could be argued that the distribution of the population is more relevant than the distribution of land.

We have not attempted to decide between an area-based criterion and a population-based criterion on theoretical grounds. Instead, we investigate both, as well as a third rule that gives each criteria equal weight. Our area-based rule assigns each station to the target unit where it has the most land. Our population-based rules assign each station to the unit where it has the most population. The population-and area-based rule weights each criterion equally. Let $PCAREA_{su}$ be the percentage of the land area of police station s that falls within target unit u. Similarly, let $PCPOP_{su}$ be the percentage of the population of police station s that falls within u. Our area-based rule allocates police station s to the u with the highest value for $PCAREA_{su}$, our population-based rule allocates s to the u with the highest value for $PCPOP_{su}$, and our population-and-area-based rule allocates s to the u with the highest value for $PCPOP_{su} + PCAREA_{su}$. We estimate percentages of land area and population using the meshblock-based dataset described in the Data section above.

There remains the question of which population measure to use. For most purposes, it makes sense to use a 'usually resident' measure, that is, a definition that allocates people to the place where they live. When exploring crime trends, a case can be made for using a 'workplace address' measure, that is, a definition that allocates (employed) people to the place where they work. This gives more sensible results when, for instance, calculating crime rates for central city areas. For the purposes of constructing geographical classifications, we have opted for the usually resident, on the grounds that it is more standard, and that, unlike workplace address measures, it does not rely on the availability of census data. However, it would be perfectly legitimate for a user of one of our classifications to calculate crime rates based on workplace addresses rather than usual residence.

Our normal area and population rules cannot be applied to the Auckland Motorways police station, which is not assigned any meshblocks in the Statistics NZ geographical classification. Instead, we allocate the Auckland Motorways station to the target unit containing the longest section of motorway.

3.3 Evaluating the allocation rules

Offence statistics are only sensitive to the choice of allocation rule if (i) there are substantial numbers of police stations that fall within multiple units of the target geography, and (ii) police stations that do fall within multiple units are treated differently by different rules. We calculate, for each of the target geographies, the extent to which these situations arise.

Assigning a police station to a single unit of the target geography when the station actually falls within multiple units leads to some misallocation of population. For instance, assigning Hornby police station to the Christchurch city territorial authority leads to some people who actually live within Selwyn district territorial authority being treated as residents of Christchurch city. Similarly, assigning a police station to a single unit leads to some misallocation of land. Calculating the extent of misallocation of population and land area under each rule therefore helps us to choose between the three rules.

Calculating the extent of misallocation also gives an indication of the amount of accuracy that we have sacrificed by using rules that assign stations to single, rather than multiple, units from the target geography. The rationale for splitting stations across multiple units is to better approximate these units. If assigning stations to single units already gives a good approximation, then this rationale is weakened.

We calculate the extent of misallocation using the meshblock data described in paragraph 3.1. For instance, we identify all meshblocks that have been assigned to the wrong territorial authority under a particular allocation rule, and then add up the populations and land areas of these meshblocks.

The three allocation rules all guarantee that every police station will be assigned to a unit of the target geography. However, they do not guarantee that every unit of the target

geography will receive a police station. Target units that are not allocated police stations can be dealt with on a case by case basis. However, it is less arbitrary and more convenient if we can avoid these sorts of retrospective adjustments. We calculate how many target units are not allocated police stations for each combination of rule and target geography.

We also calculate offence rates for one particular type of offence, serious assault, under each combination of rule and target geography, to see how much difference the rules make in practice. The reason we use serious assault is that there is huge public interest in violent crime, and the most common type of violent crime is serious assault. To facilitate comparisons across hundreds of numbers, we use multiple-panel graphics constructed using the lattice package from the programming language R (R Development Core Team 2011; Sarkar 2008).

4　Results

4.1 Proportion of police stations belonging to multiple units of target geography

The relationship between police stations and target geographies is summarised in table 2. Most police stations fall neatly within regional councils. Of the 288 police stations in the country, only 26 cross a regional council boundary. Moreover, when police stations do not fall within a regional council area, choosing which unit to put the station into is usually easy, because the population and area criteria lead to the same allocation. For instance, in 25 of the 26 cases where a police station crosses a regional council boundary, the population and area criteria produce the same allocation.

Table 2

Distribution of police stations by relationship to target geography

Relationship between police station and target geography	Target geography			
	Regional councils	Territorial authorities	Urban areas (zones separated)	Urban areas (zones combined)
Police station falls entirely within a unit of the target geography	262	146	202	202
Police station crosses boundaries of target geography				
Population and area criteria lead to same allocation	25	128	49	49
Population and area criteria lead to different allocations	1	14	37	37
Total	288	288	288	288
Source: Statistics New Zealand				

The relationship between police stations and territorial authorities or urban areas is less tidy. As can be seen in the second column of table 2, almost 50 percent of police stations cross territorial authority boundaries, and there are 14 cases where population and area criteria lead to different decisions. Fewer police stations cross urban area boundaries, because many stations lie entirely outside an urban area. However, in 37 of the 86 cases where a police station does fall within two or more urban areas, the population criterion and the area criterion pull in different directions. This typically occurs when a police station includes suburbs and surrounding countryside.

4.2 Correct allocation of population and land area

Table 3 shows the extent to which the three allocation rules assigned people and land to the correct unit of the target geography. It shows, for instance, that the population-only rule allocated 98.3 percent of the national population to the correct territorial authority, and 95.4 percent of the national land area. The highest percentage for each combination of target geography and criterion is shown in **bold**. Figure 5 illustrates what this degree of

correspondence means in practice. It compares actual territorial authorities with territorial authorities created under the 'population and area' rule, which, as reported in table 3, correctly allocates 96.2 percent of land area.

Table 3

Percent of population and land area correctly allocated, by target geography and allocation rule

| Target geography	Allocation rule			
	Criterion	Population only	Area only	Population and area
Regional councils	Population	**99.7**	99.4	**99.7**
	Land area	96.8	**97.4**	96.8
Territorial authorities	Population	**98.3**	89.9	96.2
	Land area	95.4	**96.9**	96.0
Urban areas (zones combined)	Population	**97.3**	76.9	92.2
	Land area	93.8	**98.9**	97.4
Urban areas (zones separated)	Population	**97.0**	76.6	91.5
	Land area	93.7	**98.9**	97.5

Source: Calculated from the meshblock data described in section 3.

Note: The highest percentages for each combination of target geography and criterion are shown in **bold**.

The results in table 3 imply that the areas we constructed from police stations generally do well at capturing the actual distribution of population and land area. The 'population only' outperforms the 'area only' rule at creating units that give the correct population totals, while the 'area only' rule does a better job at creating units that give the correct land areas. It is perhaps predictable that single-criterion rules perform well when judged by the criteria on which they are based. What is less predictable is that the hybrid 'population and area' rule performs almost as well as the corresponding single-criterion rules on both criteria. For instance, when applied to territorial authorities, the 'population and area' rule places 96.2 percent of the national population and 96.0 percent of the national land area in the correct unit.

Figures 5–7 illustrate the match between the actual and created geographies. The actual geographies are in shades of green, while the boundaries of the geographies created from combining police stations are in black.

Figure 5

Actual territorial authorities and territorial authorities created from police stations under the 'population and area' rule


Source: Statistics New Zealand

Figure 6

Actual and created Christchurch urban area

Source: Statistics New Zealand

Figure 7

Actual and created Auckland urban area

Source: Statistics New Zealand

4.3 Allocation of police stations to all target units

The allocation rules guarantee that every police station is assigned to a target unit, but do not guarantee that every target unit receives a police station. Table 4 shows what happens in practice. Every regional council receives at least one police station, regardless of the allocation rule used. In the case of territorial authorities, only the residual category 'area outside territorial authority' fails to receive a police station under the 'population only' and 'population and area' rules, but three territorial authorities miss out under the 'area only' rule. The number of urban areas without police stations is greater again, except under the population only rule. Non-assignment is thus a problem for urban areas, particularly when the 'area only' rule is used.

Table 4

Target areas not allocated any police stations

Target unit	Allocation rule		
	Population only	Area only	Population and area
Regional councils	[None]	[None]	[None]
Territorial authorities	Area Outside Territorial Authority	Papakura district Hamilton city Napier city	Area outside territorial authority
Main urban areas (zones separated)	[None]	Whangarei Cambridge zone Te Awamutu zone Rotorua Gisborne Napier zone Wanganui Upper Hutt zone Invercargill Kapiti	Whangarei Cambridge zone Te Awamutu zone Rotorua Wanganui
Main urban areas (zones combined)	[None]	Whangarei Rotorua Gisborne Wanganui Invercargill Kapiti	Whangarei Rotorua Wanganui
Source: Statistics New Zealand			

# 4.4 Sensitivity of offence rates to choice of allocation rule

Next we assess the sensitivity of reported serious offence rates to the choice of allocation rule. Figure 8 displays crime rates by year by regional council, for each of the three rules. The fact that, for regional councils, the plotting symbols almost all overlap means that the choice of allocation rule has virtually no effect on estimated crime rates at this level. The only visible difference is for Northland, but even here the effect is trivial compared to the variation in rates over time or between regions. In contrast, for territorial authorities, shown in figure 9, the choice of rule does matter. Trends in Invercargill, Southland, and Waipa are all appreciably different under the 'area only' rule compared with the other two rules. Different rules also yield slightly different results in places such Hamilton, Otorohanga, Waikato, and Western Bay of Plenty.

Choice of rule has a dramatic effect on some of the zone-disaggregated urban areas shown in figure 10. New Plymouth has the highest crime rates in the country under the 'area only' rule, and moderate rates under the other two rules. Nelson has falling rates under the 'area only' rule, and rising rates under the other two. For the larger zone-urban areas, however, the choice of rule mostly has a small effect on reported crime rates.

Results for the zones-combined urban areas, shown in figure 11, parallel those for the zones-separated areas. The choice of rule sometimes leads to substantially different results for small areas, but not for larger ones.

Figure 8

Serious assaults, by regional council, under three allocation rules


Source: Statistics New Zealand

Figure 9

Serious assaults, by territorial authority, under three allocation rules

Figure 10

Serious assaults, by urban area (zones separated), under three allocation rules

Source: Statistics New Zealand

Figure 11

Serious assaults, by main urban area (zones combined), under three allocation rules

5   Discussion

We have examined three rules for assembling police stations into geographical units that approximate regional councils, territorial authorities, and urban areas. Of the three rules, the 'population and area' rule performs best. Geographical units constructed with this rule closely approximate the true units, as measured by population and land area. In addition, all units are allocated at least one police station, with the exception of residual categories and small urban areas.

We have restricted our comparisons to rules that assign police stations to single units from the target geography. Our main reason for doing so is to maximise simplicity and transparency. Assigning stations to multiple units might lead to more accurate crime statistics by more faithfully recreating the target geographical units. However, the 'population and area' rule seems to work sufficiently well that the scope for further improvement is small. For instance, 96 percent of the population is assigned to the correct territorial authority. Rules that assign police stations to single units seem to offer a better trade-off between simplicity and accuracy.

Further reassurance is provided by the case study of serious assault statistics. These statistics are not, in general, sensitive to the choice of allocation rule. The main exception is small urban areas: different allocation rules do, in fact, lead to different trends or levels for serious assault rates. Small urban areas aside, the lack of sensitivity is reassuring because it means that people who disagree on the correct choice of rule can nevertheless agree on regional crime trends.

Overall, our results suggest that it would be feasible to create geographical crime statistics, and that a good way to do so would be to apply the 'population and area' rule. These statistics could be created for regional councils and territorial authorities, and for large urban areas, though not for small urban areas. To say that statistical series are feasible is not to say that they should or will be produced by Statistics NZ. This depends on factors such as stakeholders' priorities, on the availability of resources, and on Statistics NZ's wider strategy for administrative data.

This does not mean that the statistics are error-free. For instance, the trend towards placing police officers in centralised police stations to service surrounding smaller surrounding stations may be reducing the importance of police station boundaries, and hence the accuracy of the station-level statistics[1]. Whenever interpreting offence statistics, it is always necessary to consider the possibility that an apparent trend or differential may be the result of reporting patterns. When making such judgements, there is no substitute for detailed knowledge of the process used to produce the statistics.

---

[1] We are grateful to an anonymous referee for describing this phenomenon.

6 References

Freedman, M, Lane, J, & Roemer, M (2008). New approaches to creating data for economic geographers. *Journal of Official Statistics* 24(1): 133–156.

Gregory, I N, Marti-Henneberg, J, & Tapiador, F J (2010). Modelling long-term pan-European population change from 1870 to 2000 by using geographical information systems*. Journal of the Royal Statistical Society: Series A*. 173(1): 31–50.

New Zealand Police (2008). *New Zealand Police National Recording Standard*. Wellington: Author.

R Development Core Team (2011). *R: A language and environment for statistical computing.* Vienna: R Foundation for Statistical Computing.

Rees, P (1986). Choices in the construction or regional population projections. R Woods and P Rees (eds.) *Population Structure and Models: Developments in Spatial Demography*. London: Allen and Unwin.

Sarkar, D (2008). *Lattice: Multivariate Data Visualization with R*. Springer: New York.