

Estimating bird abundance: making methods work

STEPHEN T. BUCKLAND, STUART J. MARSDEN and RHYS E. GREEN

Summary

In many bird monitoring surveys, no attempt is made to estimate bird densities or abundance. Instead, counts of one form or another are made, and these are assumed to correlate with bird density. Unless complete counts on sample plots are feasible, this approach can easily lead to false conclusions, because detectability of birds varies by species, habitat, observer and many other factors. Trends in time of counts often reflect trends in detectability, rather than trends in abundance. Conclusions are further compromised when surveys are conducted at unrepresentative sites. We consider how to avoid these problems. We give a brief description of distance sampling methods, which allow detectability to be estimated. We consider strategies to ease their implementation, to enhance their reliability, to adapt the methods for difficult species, and to deal with circumstances in which representative sampling is problematic. We also consider some of the common problems encountered, and suggest solutions.

Introduction

Estimates of bird abundance (population size or density) are widely used in bird conservation. For example, they allow us to measure changes in population size and hence gauge the impact of habitat loss, pollution or harvesting (e.g. Lambert 1993) and to assess whether or not isolated populations are viable (e.g. Githiru and Lens 2006). They are also one of the cornerstones of the IUCN Red List classification scheme (IUCN 2001). It may seem obvious that abundance estimates should have low bias (minor systematic under- or overestimation) and high precision (a low level of uncertainty). However, in practice this can be difficult and expensive to achieve. As a consequence, numbers of birds recorded during some standardised fieldwork are often reported as an index of abundance and the argument used that estimates of absolute density or population size, rather than relative measures, were not necessary. However, such counts are frequently used to infer a change over time in abundance or a difference between regions or habitats. The conclusions reached may trigger conservation action to remedy a perceived population decline or to assign higher priority to some habitats or areas than others for protection or conservation management. Conservation resources may be misapplied if the inferences about trends or differences in abundance are inaccurate. Two frequent causes of inaccuracy are, first, that the proportion of birds present in the survey area that are recorded varies according to time and place (i.e. detectability is variable) and, second, that the places selected for surveys are not representative of the whole area for which the conclusions are drawn.

There are three solutions to the problem of variation in detectability. The first is to design the fieldwork so that all birds within the surveyed areas can reasonably be assumed to be detected. For example, you may have plots that are sufficiently small to ensure complete enumeration of birds on them. The difficulties for this approach are, first, ensuring that you do enumerate all individuals within each plot, especially for multi-species surveys; second that bird movement into, out of, and around the plots during the count can lead to biased estimates; third that, if plots are small enough to allow reliable complete enumeration, many birds are likely to be detected beyond the boundaries of plots, and these data cannot be used, so that the method is inefficient (particularly frustrating with rare species); and fourth, in order to detect all birds within the plot,

the observer may cause greater disturbance, so that birds may leave the plot without being detected.

The second solution is to design the survey and field procedures so that detectability can be assumed constant across the units of time (e.g. between years) or space (e.g. across habitats/land uses) that you are comparing. This can be difficult to achieve if, for example, vegetation conceals birds to different degrees at different times or places, if observers change, or, for monitoring surveys, if there is habitat degradation or succession, or trends over time in factors that affect detectability such as traffic noise, disturbance and observer experience and hearing ability.

Because the first two solutions are often difficult to achieve in practice, we concentrate in this paper on the third solution, in which detectability is estimated, so that abundance can be estimated even though we do not have complete enumeration of individuals on the sample plots. The main disadvantage of this approach is that it adds complexity, and more information typically needs to be recorded. We argue here that these issues are not generally a serious impediment to the strategy. The availability of user-friendly software (Distance, Thomas *et al.* 2006) to analyse the data, and an increasingly large pool of ecologists across the world who are experienced in its use, go a long way towards addressing the issue of additional complexity in the analysis.

The problem of unrepresentativeness of the places selected for surveys applies regardless of the methods used to estimate density. There is much useful information on sampling and study design in various texts (e.g. Bibby *et al.* 2000, Greenwood & Robinson 2006) but we suggest later some ways to avoid the problem of unrepresentativeness, or to allow for its most serious effects, in a distance sampling context.

Distance sampling

Distance sampling (Buckland *et al.* 2001) is an extension of plot sampling, in which birds are counted within a sample of defined areas (plots). The plots are long, narrow strips (line transect sampling) or circles (point transect sampling). Simple plot sampling assumes that all of the birds present in the plot just before the arrival of the observer are counted without error. The distance sampling extension allows for the possibility that some of the birds present in the plot are not detected. However, it is useful to begin by considering the case of perfect detection.

Fixed-radius point counts

In this method, an observer goes to each of a number of points within the survey region of interest. These points should be randomly placed, or more commonly, on a systematic grid randomly positioned over the survey region. All detected birds within a fixed radius are counted. If the radius is small, then it may be reasonable to assume, at least for a subset of species, that all birds on the plot are detected. In that case, we can estimate bird density by

$$\hat{D} = \frac{n}{a} = \frac{n}{k\pi w^2} \quad (1)$$

where k is the number of points in the design, n is the number of birds counted (summed across the k plots) of the species of interest, w is the fixed radius of the plot, and $a = k\pi w^2$ is the total size of the surveyed plots, termed the 'covered area'. Abundance in the wider survey region of interest is then estimated by multiplying this density by the size A of the survey region: $\hat{N} = A\hat{D}$.

Strip transect sampling

If, instead of counting from points, the observer travels along lines, counting all birds within a fixed distance w of the line, then this is strip transect sampling and w is called the strip

half-width. The lines should be positioned according to some randomised design; this is generally a systematic grid of parallel lines, randomly superimposed on the survey region. Lines may traverse the entire region (in which case they will typically vary in length), or may be short segments, systematically spaced (Strindberg *et al.* 2004). Density is again estimated by dividing the count n by the covered area, $a=2wL$, where L is the total length of transects. Thus

$$\hat{D} = \frac{n}{a} = \frac{n}{2wL} \tag{2}$$

and $\hat{N} = A\hat{D}$ as before.

Point transect sampling

Now suppose that an unknown proportion $1-P_a$ of birds within the circular plots of radius w are undetected. If we can obtain an estimate \hat{P}_a of the proportion detected, then we can estimate density by

$$\hat{D} = \frac{n}{a\hat{P}_a} = \frac{n}{k\pi w^2\hat{P}_a} \tag{3}$$

In standard point transect sampling, we estimate P_a by modelling the detection function $g(r)$ say, which is the probability that a bird at distance r from the point is detected, with $0 \leq r \leq w$. We assume that $g(0)=1$ (that is, any bird at the point is certain to be detected), and choose a suitable model for this function. This choice specifies the shape of the relationship between the probability of detection and the distance between the bird and the point. P_a is the average probability of detection of birds in the covered area, where we average over distance r from the point. Mathematically, this may be expressed as $P_a = \frac{2}{w^2} \int_0^w rg(r)dr$ (Buckland *et al.* 2001:54). We use maximum likelihood to estimate the parameters of $g(r)$ (Buckland *et al.* 2001:61–68), which allows us to estimate P_a and hence density D and abundance $N=A\hat{D}$.

If a detection event can comprise more than one bird, we need also to estimate mean group (termed ‘cluster’) size in the population. In this case, equation (3) gives estimated cluster density, and we need to multiply this by our estimate of mean cluster size to estimate bird density. The mean size of detected clusters tends to be a biased estimate of mean cluster size in the population, because bigger clusters tend to be more detectable, so other methods are used (Buckland *et al.* 2001:71–76).

Line transect sampling

Again, we suppose that an unknown proportion $1-P_a$ of birds within the covered area are undetected. Then

$$\hat{D} = \frac{n}{a\hat{P}_a} = \frac{n}{2wL\hat{P}_a} \tag{4}$$

In standard line transect sampling, we specify a model for the detection function $g(x)$ say, which is now the probability that a bird at distance x from the line is detected, with $0 \leq x \leq w$. Note that we do not usually distinguish whether the bird is to the right or to the left of the line. P_a is the average detection probability, averaging over distance from the line, for birds that are within distance w of a line, giving $P_a = \frac{1}{w} \int_0^w g(x)dx$ (Buckland *et al.* 2001:53). As for point transect

sampling, we use maximum likelihood to estimate the parameters of $g(x)$ and hence P_a , from which density and abundance are estimated as before.

If birds occur in clusters, we need to multiply equation (4) by our estimate of mean cluster size in the population to estimate bird density.

Assumptions of distance sampling methods

We make the following important assumptions. In this section, we just note them, and comment on the effects if they are violated. We will later consider ways to meet the assumptions, relax them, or make our methods more robust to their failure.

(a) *Birds on the line or point are certain to be detected.* Birds may remain concealed and quiet and hence remain undetected even when they are at the point or on the line. If we fail to detect 40% of birds on the line or point, we expect a negative bias in our abundance estimate of around 40%.

(b) *Birds are detected at their initial location.* Distance sampling methods are ‘snapshot’ methods; conceptually, we ‘freeze’ the birds in position. If birds move in response to the observer before detection, either towards or away from the line or point, then considerable bias can result. If movement is independent of the observer, bias is slight for line transect sampling provided on average the birds are moving at less than half the speed of the observer (Buckland *et al.* 2001:131). For point transect sampling in which birds are recorded over a period of time, movement of birds through the plot will lead to the overestimation of abundance because birds not initially present in the plot enter it. The longer the period of recording, the larger this effect is likely to become. Even random movements of birds within the plot lead to overestimation of abundance if they result in individuals being double counted, and because the bird is more likely to be detected if it moves closer to the point, giving a downward bias in the recorded distances. Bibby *et al.* (2000:96) comment that some studies have quoted density estimates up to ten times reality due to these effects.

(c) *Distance measurements are exact.* If for example we overestimate all distances by 10%, then line transect estimates of abundance are biased low by about 9%, and point transect estimates are biased low by about 17%. If we underestimate all distances by 10%, then line transect estimates are biased high by about 11%, and point transect estimates by about 23% (Buckland *et al.* 2001:132,176). If we have errors in our estimated distances but the mean error is zero (i.e. the estimates are unbiased), line transect estimates are usually approximately unbiased, whereas point transect estimates tend to be biased high. The greater the error variance of the estimates, the greater is the bias. Buckland *et al.* (2001:195–196) give a point transect example in which the coefficient of variation of the estimated distances was 35% (that is, the standard deviation for the estimated distance was 35% of the true distance), which resulted in positive bias in estimated density of 85%.

(d) *For birds that occur in clusters (groups), cluster sizes are recorded without error.* In some circumstances, cluster sizes may be accurately estimated close to the line or point, but poorly estimated (usually underestimated) at larger distances. Bias from this source can be avoided by using the regression correction for size-biased sampling that is the default for clustered data in the software Distance (Buckland *et al.* 2001:73–75). Note that Distance sometimes gives a warning that the size-bias correction is in the wrong direction in this case, because underestimation of size for more distant clusters is greater than the upward bias arising because larger clusters are more detectable. This warning should be ignored when it is suspected that the size of distant clusters is underestimated.

(e) *The sampled plots (circles or strips) are representative of the entire survey region.* This is not usually listed as an assumption, because if an appropriately randomised design is used, the

assumption holds by design. However, non-random plots are often covered, in which case this assumption becomes important (see below).

Although other assumptions are made, generally only the above five have any practical significance.

Strategies to make the methods easier to apply

Defining the area and population under study

Distance sampling methods have now been used in a wide range of habitats and regions, and at a broad range of spatial scales. At one end of the spectrum are schemes such as the UK Breeding Bird Survey (BBS, Freeman *et al.* 2007) which aim to determine population sizes and trends in birds at a countrywide scale (Newson *et al.* 2005, in press). In such cases, the populations and the region are well-defined, and large numbers of fieldworkers are required to ensure representative coverage.

In many large-scale surveys, the population and area are much less clearly defined. For example, the UK Common Birds Census (CBC), now discontinued, was based on non-random sites mostly in farmland or woodland, so that the populations were loosely defined as the UK farmland and woodland populations of each species, although parts of the UK were often excluded from inference due to inadequate survey coverage (Gregory *et al.* 2000). Even within the surveyed habitats, sampling was very non-random, so that densities on surveyed plots were likely to be biased estimates of average densities within farmland and woodland. Freeman *et al.* (2007) discuss the implications of these shortcomings for monitoring UK bird populations. By ensuring representative sampling of sites, the UK BBS avoids the complications of interpretation that analysts of CBC must deal with.

Distance sampling is also employed at much smaller scales, for example to determine population sizes of species within individual protected areas (Marsden *et al.* 2005) or habitat fragments (Buckland 2006; Mulwa *et al.* 2007), or to determine density differences across habitats within a defined study site (Marsden *et al.* 2006). In these cases, we must restrict our inferences concerning abundance to within the study site, as abundance or abundance shifts might not reflect those across the wider landscape. For example, bird densities may be lower towards the edge of a species' range or they may be influenced by different harvest pressures, while a species may be declining in some parts of its range at a time when densities are stable or even increasing in others (e.g. Conti Nunes *et al.* 2007). In some studies, the transects covered are not spread throughout the study site, or are too few to provide adequate replication, so that inference should be further restricted to just the strips surveyed, or possibly a subset of the study site which has been representatively sampled (e.g. Cabanne *et al.* 2007).

Which survey method?

A fundamental consideration for distance sampling surveys of birds is whether to conduct a line transect or a point transect survey. Line transect surveys have the following advantages: for the same effort, they generate more detections, because time is lost in moving between points for point transect sampling; given the same number of detections and good quality data, line transect estimates tend to have lower bias and higher precision, essentially because a higher proportion of detections is close to the observer than for point transects. Hence point transect sampling should be considered only if it is believed that the quality of line transect data might be severely compromised. This may occur because of the impossibility of navigating a random transect line in dense habitats or difficult terrain – it is easier to get to a random point than to travel along a random line. It may also occur in multi-species surveys, where the observer risks being 'swamped' with bird records. It is easier to concentrate on detecting and recording everything if the observer doesn't also have to navigate along a line. This is often a major consideration (both

in terms of data quality and safety) in surveys within tropical forests where terrain is difficult and birds occur at all vegetative layers from the ground to the forest canopy. There is a practical advantage for the analyst of using point transects: it is easier to record habitat at a point than along a line, making it simpler to investigate habitat associations. It is also easier to model variation in density through the survey region using point transect data (Hedley *et al.* 2004). Line transects can be divided into short sections and habitat information collected for these, although analysis is less straightforward due to correlation in the data from successive line segments.

Another consequence of an observer being 'swamped' when conducting a multi-species survey is that it may be impossible to estimate accurately the distance to every detected bird, especially when most detections are aural. Bibby and Buckland (1987) and Buckland (1987) developed binomial point transect models with two distance categories, allowing observers to record whether or not each bird is within some radius r_o . The observer then simply needs to tally at each point the number of birds detected in two categories: 'near' and 'far'. For woodland habitats, r_o might be of the order of 25 or 30m. Only those birds at a distance close to r_o need to be located accurately so as to avoid misclassification.

The above two-category models assume that data are untruncated; that is, w is assumed to be infinite. This means that observers do not have to make any judgements of whether distant birds are within or beyond w . The software Distance (Thomas *et al.* 2006) uses numerical methods for which a value for w must be specified, if distances are collected in distance categories as here. By selecting a value for w that is beyond the largest distance at which a bird is likely to be detected, a good approximation to an untruncated two-category model can be obtained.

For the UK Breeding Bird Survey (Freeman *et al.* 2007), line transect sampling is conducted, and observers record whether birds are in one of three distance intervals: 0–25m, 25–100m and >100m. In analysis, data from the third category are generally not used; i.e. the data are truncated at $w = 100$ m. Newson *et al.* (2005, in press) fitted truncated two-category models with $r_o = 25$ m and $w = 100$ m using the Distance software, to estimate the size of national populations.

A problem with two-category (or binomial) data is that they provide too little information on the shape of the detection function to allow goodness-of-fit testing, so that a model must be uncritically assumed. To avoid this, more distance categories would be needed, perhaps four or five. (If it is feasible to have more categories, then it is probably preferable to record actual distances rather than use grouped data.)

When to survey

Distance sampling is most efficient when birds are spread relatively evenly through the survey region. Thus for a resident songbird that forms territories during the breeding season and small flocks at other times, surveys during the breeding season should give better precision. By contrast, complete count methods are more workable if birds are in a few large groups, provided that all groups can be located. Counts of non-breeding wader and wildfowl flocks and of breeding colonies are clear-cut examples of this strategy, but other cases, such as species that occur in clusters, leks or loose colonies within say tropical forests (e.g. *Picathartidae*, *Pipridae*) pose the problem of when to switch from a distance sampling to a complete count or plot sampling approach. In such situations, and especially with poorly-known species, pilot studies are essential to determine likely distribution and encounter rates. Adaptive distance sampling might be a useful approach in some situations (Buckland *et al.* 2004).

Care is needed in deciding the timing of surveys of tropical forest birds. Differences in calling/singing rates, home range size, short-term movements in response to shifts in resources, differential habitat use, and participation in large mixed-species foraging flocks (e.g. Maldonado-Coelho and Marini 2004) can all affect detectability, which may feed through to effects on precision and bias of the resulting abundance estimates.

Strategies to enhance the reliability of the methods

Snapshot methods

As noted above, distance sampling (in common with other plot sampling methods) is a 'snapshot' method, and is therefore biased in the presence of bird movement, because surveying cannot be done instantaneously. This can be especially problematic for point transect sampling, in which the observer stays at a single location for several minutes. Bibby *et al.* (2000:97) recommend five-minute counts in temperate regions, while Lee and Marsden (2008) suggest count periods ranging from four to 10 minutes dependent on the type of forest bird surveyed. In many surveys, longer counts are conducted, perhaps up to 20 minutes, to enhance detectability. This may be necessary for example for songbird species, if the gap between songbursts often exceeds five minutes, to ensure certain or near-certain detection of birds at the point. However, for species that typically move around within their territories during the period of a count, two problems arise, both of which cause overestimation of density. First, some birds not initially present in the plot move into it during the recording period, and second, birds within the plot are more likely to be detected when close to the point than when further away, leading to too many short distances (see earlier). Both of these problems may be circumvented by using a snapshot approach (Buckland *et al.* 2001:173). This involves defining a moment (say two minutes after arrival at the point), at which a snapshot of distances from the point is recorded. The observer uses the time before the snapshot moment to detect and locate birds and, as far as possible, keep track of any movements. The positions of those birds are then recorded at the snapshot moment. The observer uses as much time as is needed to confirm locations of birds after the snapshot moment. He or she may also move away from the point at this time, to help triangulate for aural detections. In practice, the presence and location of many individual birds is inferred, rather than being known, at the snapshot moment. For example, a bird which cannot be seen but which sings at the same location both before and after the snapshot moment is inferred to have been there at that moment. Birds whose location is unknown, and cannot be reliably inferred, at the snapshot moment are not recorded; that is, they are treated the same as undetected birds. Provided all birds at or very near the point can be located, this does not bias the method: the 'detection function' under this approach is actually the probability that a bird can be located at the snapshot moment as a function of its distance from the point, rather than its probability of detection. The method was successfully tested on four species by Buckland (2006), giving on average lower estimated bias and higher precision than five-minute counts.

Estimating distances

Inability to estimate distances accurately is often cited as a reason for not using distance sampling. We noted above that data may be grouped into distance intervals to reduce this problem; in the extreme case, observers simply record whether birds are within or beyond an easily measured or judged distance r_0 . For surveys with properly trained observers, it is generally possible to record ungrouped distance data. Laser rangefinders have made this task very much easier for terrestrial surveys, giving accuracy to the nearest metre over distances of several hundred metres. They are also relatively inexpensive, and should be considered as standard equipment for terrestrial surveys of birds that use trained observers. Even when most detections are aural, it is often possible to locate the tree or bush in which a bird is located, or to estimate that the bird is at approximately the same distance as a visible object. Accurate measurement of distance to that estimated location removes one source of error.

In dense woodland, accurate distance estimation can be a severe limitation (Alldredge *et al.* 2007b). In this circumstance, observers should be trained by estimating distances to birds whose location has been determined by other observers. Field methods should be explored to allow observers to check and improve upon their estimates. An assessment should be made of accuracy

and precision as a function of distance from the point, as these often deteriorate quickly beyond a certain distance. If the quality drops off too badly, data beyond this point might be truncated. Ignoring bird records at these large distances from the observer may have an added advantage that recorders can better concentrate on birds much closer to them, thus better ensuring that detection probabilities very close to the transect line or point are as high as possible. Estimation of distance to birds in tropical forests may be particularly problematic, given the vegetation complexity and height, and the fact that usually > 80% of bird contacts are heard only (e.g. Marsden *et al.* 2006). Rigorous training is essential in such cases and it may be wise for recorders to concentrate wholly on bird detections within 30 m or so of the transect line or point.

In line transect sampling, an observer is moving, which allows triangulation to estimate distances to aural cues. In point transect sampling, if there is a rigid protocol that requires the observer to remain at the point, distance estimation for aural cues may be compromised. If the observer is allowed to move around within limits, distance estimates may be improved substantially. If such movement may disturb birds or lead to failure to detect all birds at or near the point, a possible strategy is to require the observer to remain at the point for say three minutes, then allow movement to aid distance estimation in the last two minutes of a count. Another solution to the problem of estimation of distances to aural cues is to have two (or more) observers, a few metres apart, again allowing triangulation. Multiple observers can also be used to enhance reliability, by assigning a set of species to each observer. This is especially useful if some species require a different field protocol than others, to yield reliable estimates. Increasingly, observers work in pairs for safety reasons, and this offers options for improved field methods.

Although not available yet, it seems likely that acoustic arrays will soon be developed that will allow accurate estimation of distance to singing or calling birds, at least for some species. If this can be coupled with automated identification of songs and calls, the method has considerable potential (Brandes 2008).

Which age classes and sexes of birds should be recorded?

In the UK Breeding Bird Survey, observers record all adult birds, because this keeps fieldwork simple when all detected species must be recorded. However, for most territorial songbirds, more reliable estimates are likely to be obtained by recording singing males only, assuming that this gives a good guide to the number of territories. While territory-holding males are easy to detect in early morning in late spring and early summer in temperate climes, females may well violate the assumption of certain detection on the line, especially if some are on nests, and their detectability is likely to fall off much more quickly than for singing males. Data pooled across the sexes are therefore less reliably modelled than data on singing males alone. There are many examples of species for which this may also hold in tropical forests and other habitats. Similar considerations apply to the recording of immature, as well as adult birds, where the behaviour that affects the detection of the birds varies among age classes. There is a trade-off here between including as many individual birds as possible, thereby increasing the sample size for the estimation of the detection function, and including classes of birds which have widely differing detection properties, leading to unreliable modelling of the detection function.

Birds in flight

Birds detected in flight during point counts or line transects present a problem because they are moving much faster than an observer on foot, so generate substantial upward bias in density estimates, if no special measures are taken. For point transect sampling, the snapshot method resolves the problem: birds in flight are only noted if they are visible (or audible) at the snapshot moment, and their distance from the point at that moment is estimated.

Another option is to estimate density for birds not in flight, and then separately estimate the average proportion of time that a bird is in flight, from which total abundance may be estimated (Marsden 1999). This strategy may also be used for line transect sampling (Buckland *et al.* 2001:202). In this case, there is another option for flying birds that typically cross the transect just once while the observer is recording (as distinct from species that mill around): record the distance of the bird from the line when it passes 'abeam' of the observer (*i.e.* when its direction from the observer is at right-angles to the line; Buckland *et al.* 2001:202). If the bird is no longer detectable at this time, then it is not recorded. Provided that the birds are moving independently of the observer, resulting density estimates will be approximately unbiased. This method might be especially suited to shipboard line transect surveys of seabirds. For most large forest birds (e.g. parrots, pigeons, hornbills or toucans) where only a small proportion of the population is in flight at any one time, little bias results from simply excluding such detections from the analysis. Some raptors may be an exception.

Strategies for difficult species

Cue count methods

For species that are primarily detected by a cue (typically a songburst or call), and that tend to move around during the typical duration of a point count, cue count methods should be considered. Cue counts were developed for shipboard transect surveys of whales, for which a cue was a whale blow (Hiby 1985). The idea is that the distance from the observer to each detected cue within a sector ahead of the ship is recorded, and these are analysed in much the same way as point transect data to estimate the number of cues per unit time per unit area. The cue rate is estimated from separate fieldwork: the average number of blows per unit time per animal. This allows cues per unit time per unit area to be converted to an estimate of the number of animals per unit area (Buckland *et al.* 2001:191–197).

Cue counts can be conducted along line transects or from points, and the same analysis applies in either case (Buckland *et al.* 2004:356–359). The measure of effort used in the calculation of abundance is time spent in the field recording, rather than line length or number of visits to points. In the case of point transects, because movement independent of the observer does not bias the method, longer can be spent at each point. This offers potential for acoustic arrays, which could be left operating for several hours. Automated distance estimation and species identification would make such an approach very attractive. Estimation of the rate at which individual birds produce cues should be conducted at the same time as the main survey, to ensure that the estimated cue rate is representative of the rates that applied in the main survey. Vocalisation rates for a species might be expected to vary by sex and age of individual, by time of day, by season, by habitat, and under different weather conditions, and may be influenced by population density itself (van Dongen 2006, Staicer *et al.* 2006). Consequently, estimation of cue rate should be conducted in a well-designed survey of sufficient size.

Cue counts have the advantages that detection of a bird at or near the point can be uncertain, provided it is certain to be detected if it gives a cue; movement of the bird independent of the observer does not generate bias; and there is no need to identify whether several different cues detected in the same plot were made by the same bird. Disadvantages are that the cue rate per bird must be estimated by monitoring a representative sample of birds; observers are more likely to be swamped, as they need to estimate a distance to each individual cue, rather than each bird; and multiple cues from one bird generate non-independent distance data that are markedly 'overdispersed'. The last disadvantage makes selection among various plausible forms of the detection function difficult, as pseudoreplication makes goodness-of-fit tests and model selection criteria such as Akaike's Information Criterion (AIC) invalid.

Buckland (2006) trialled cue counts from points for four species of songbird, and concluded that the method was workable for specialist surveys of single species, but that it would not be workable for large multi-species surveys.

Lure point transects

For some species, very few records are obtained because they are rare, because they easily pass unnoticed even if they are right on the line or point, or because the detection probability may drop so quickly that sample size for standard distance sampling is inadequate. If birds can be attracted in to a point, then a larger area can be effectively surveyed. This is the idea behind lure point transects (Buckland *et al.* 2006). The 'lure' might be taped calls or song, which should be of a standardised type, volume and duration and played in a standard way, for example by directing the speaker for an equal time in each of four directions at right angles to one another. The problem with the method is that we do not know the initial positions of birds (i.e. how far they were from the point), before they were attracted by the lure. A series of trials is therefore conducted. For each trial, a bird or group of birds is located by an observer. He or she then contacts a second observer, usually by radio or telephone, who plays the lure in the standard way at a predetermined distance from the birds, and records whether or not there is a response. If birds do respond, the first observer can confirm whether it was the birds he or she was monitoring. Recording by both observers of the times at which events occur is helpful for this. Care should be taken that the observer who plays the lure does not allow knowledge of the location of the first observer to bias the way in which the lure is played. A detection function is estimated from the trial data. Because we know both the 'successes' (those birds that did respond) and the 'failures' (those that did not), we can model the detection function using logistic regression and distance from the lure as a covariate. (Other covariates may also be included.) This allows for uncertain detection at the point, so we do not need to assume $g(0)=1$. Buckland *et al.* (2006) used the method to survey crossbills (*Loxia* spp), but it may also be suited to surveying some ground species, such as Rallidae and kiwis (*Apteryx* spp), for which the trials may be set up using radio-tagged birds.

Removal estimators

Nichols *et al.* (2000) and Farnsworth *et al.* (2002) developed removal type estimators for estimating detectability from point counts when detection distances from the point are not recorded. The difficulty with this approach is that such estimators can be badly biased in the presence of unmodelled heterogeneity (differences among birds in the probability of being detected). The main source of unmodelled heterogeneity in this case is due to distance of the bird from the point. Farnsworth (2002) attempted to address heterogeneity by allowing two categories of bird: those with high detectability and those with low. However, heterogeneity arising from distance of the bird from the point will tend to be a smooth function of distance, with detection probability decreasing with distance. For a more detailed discussion, see Buckland *et al.* (2004:352–354). If distances are recorded, then the problem of heterogeneity may be overcome by combining removal estimators with distance sampling methods, as described by Buckland *et al.* (2004:354–356). However, in the presence of movement during the count period, the size of the plot being surveyed is not well-defined because birds initially outside the surveyed circle may enter it, so that population density and abundance in the survey region might not be reliably estimated.

Double-observer methods

Double-observer methods for line and point transect sampling allow combination of mark-recapture and distance sampling data, so that the assumption of certain detection of birds on the

line or point is not required. Double-observer surveys complicate both field methods and analysis. Typically, both observers record birds independently of the other, so that it is important that one observer does not cue the other observer when a bird is detected. Methods are also available for one-way independence, so that observer 1 (say) is unaware of detections made by observer 2, but observer 2 is aware of detections made by observer 1. This can ease the task of matching detections made by the two observers, so that birds detected by both can be identified. When detection on the line or point is not certain, estimates of abundance are biased in the presence of unmodelled heterogeneity in detection probabilities. Hence it is important to record and model covariates that affect detectability, such as bird behaviour and habitat. Bias from unmodelled heterogeneity may be reduced by assuming 'point independence' – independence between observers only on the line or point. Laake and Borchers (2004) and Borchers *et al.* (2006) give comprehensive methods for analysing double-observer data, and these are now available in the software Distance (Thomas *et al.* 2006).

Strategies for when representative sampling is difficult

Random placement of transects

It is often argued that it is not practical to conduct surveys with randomly placed plots, points or transects when volunteer observers are involved. However, the UK Breeding Bird Survey is a line transect survey conducted by volunteers, in which a stratified random sample of 1 km squares is selected. Within each sampled square, two transects, each 1 km long, are surveyed. If a national survey, dependent on large numbers of volunteer observers, can use a randomised design, then most surveys should be able to. Of course, there have had to be some compromises. For example, if a landowner refuses to allow access, a replacement square must be drawn. Also, it is often not possible to follow exactly the nominal transects; in this circumstance, observers map the actual location of their lines, allowing squares to be classified according to degree of departure from the nominal lines. This can be expected to give some bias against open field species for example, and towards species of edge habitats. Nevertheless, it converts an otherwise unimplementable design into a successful and popular survey, with nearly 3000 squares sampled annually.

Stratification is often a useful method of reducing the cost of random sampling. In the case of the UK Breeding Bird Survey, the UK was stratified into regions, with higher sampling rates in regions with more volunteer observers. Similarly, if the survey region can be stratified according to ease of access, then lower sampling intensities can be used in the less accessible strata.

If there are parts of your survey region that simply are not surveyable, then a reduced survey region should be defined, in which random sampling (possibly using stratified and/or systematic schemes) is feasible. Inferences should then be restricted to the reduced region.

Cluster sampling can be useful, especially with point transect sampling. If it is costly in time or resources to get to a location, then it is very inefficient to survey just one point at that location. Instead, a grid or 'cluster' of points centred on that location can be defined and surveyed. Data from different points within the same cluster cannot be treated as independent. The simple way to deal with this is to treat each cluster of points as a sampling unit. To achieve this in Distance, pool the bird records from points within the same cluster, and record the effort as the number of points in the cluster.

Road or path transects

It is common for transects to be placed along roads, rivers or paths. Although sometimes it simply is not possible to have randomly positioned lines or points, more often, it is a matter of time and resources. However, it can take greater resources to verify that non-random transects yield good estimates of density than to survey random transects in the first place.

There are many reasons why bird density might not be typical along roads or paths. These include greater disturbance, presence of 'edge' habitat, and avoidance of more difficult terrain (possibly with less accessible, more pristine habitat). Even relatively narrow trails can have large effects on local abundance measures: bird density estimates along an old Philippine logging trail were around 25% lower than in adjacent forest, with striking differences for canopy frugivores (higher in the forest interior) and nectarivores (higher along the logging trail) (Lee and Marsden unpubl. data). When the effect of a road on density is restricted to the strip extending a distance w either side of the road (where w is the maximum distance at which detections are recorded), then point transect data along roads can be used to correct for this effect, if angles as well as distances are recorded, although the corrections can be sensitive to the detection function model assumed (Marques 2007). Another strategy is to collect data along lines perpendicular to roads (Marques 2007), either to verify that the effect of roads is restricted to these strips of half-width w , or to correct biased estimates of density obtained from road transects. For such an analysis to have adequate power, many of these lines might be required, in which case it may be a better use of resources to use random lines rather than road transects.

With work in tropical forests, and especially in surveys of threatened species, representative placement of survey effort might be particularly challenging. First, surveyors may be tasked to undertake surveys in areas which are extremely large in relation to the level of sampling effort possible. In these cases, surveyors have to balance the costs/benefits of surveying just a portion of the site. This strategy may be time-efficient but comes at the cost of compromising inferences over the whole of the species' range or the entire protected area. Densities of tropical birds can vary greatly through large areas, swamping the effects of for example habitat alteration within small study sites (e.g. Jones *et al.* 2003). Local habitat and physical heterogeneity is high in tropical forests making positioning of lines or points very important. Random placement of point or line transects is rare in tropical forests, and common approaches are time-efficient placement along paths (Jepson *et al.* 2001), positioning point transects 50 m on alternate sides of path-based transects (e.g. Marsden *et al.* 1997), and specially cut grids (e.g. Hilton *et al.* 2003).

Problems and solutions

I have just one transect line. How do I analyse my data?

It is surprising how often this question is asked. If you have a single transect, you will not be able to make valid estimates of the standard error and confidence limits of your estimates of density and population size because of the lack of replication. If your transect is long and randomly positioned, it may record a large number of birds from a representative part of your study area and yield an accurate estimate of population size. However, this will not be of much use because you will not know how precise the estimate is and will not be able to test whether it is different from estimates for other areas or the same area at other times. To avoid this difficulty, replicate lines located according to a randomized scheme should be used. Buckland *et al.* (2001:232) recommend at least 10–20 lines.

I have walked each of my five transects ten times. Can I analyse this as 50 transects?

As noted above, we require that the transects are representative of the wider survey region. In this case, there are just five transects, which is unlikely to be adequate replication. To analyse the data as if there were 50 transects is termed 'pseudoreplication', and leads to substantial underestimation of variance.

If the reason for the repeat surveys is purely to increase the number of detections, then the way to analyse the survey is to pool the detections from the repeat visits for each line, and then record effort for that line as its length multiplied by the number of times it was surveyed – ten in the above case.

My survey region is too small to allow the recommended number of lines/points.

Distance sampling is especially efficient relative to methods such as complete count and mark-recapture when the survey region is large. For small survey regions, distance sampling is generally inefficient – in extreme cases, you may have to detect each bird in your population several times, before the precision of your abundance estimate is adequate. Hence if it is feasible to conduct a complete count, that would be preferable. If it is not, then distance sampling can still be made to work. Substantial overlap between plots can occur without appreciable bias. Buckland (2006) demonstrated this by simulation, in which most birds in the survey region were detected from at least two points, and many from four or more points. Resulting bias in density estimates was estimated as half of 1%, with upward bias in the standard error of < 9%.

For small survey regions, edge effects can be problematic. These may be alleviated by sampling into a buffer zone around the survey region, as described for point transect sampling by Strindberg *et al.* (2004:192–194) and for line transect sampling by Strindberg *et al.* (2004:200–201). Essentially, line segments or points that are outside the survey region, but within a distance w of the boundary, are surveyed, and any birds detected from them that are inside the survey region are recorded. Effort corresponding to these line segments or points is not recorded.

I do not have enough data to estimate the detection function separately for different habitats, regions or years.

Often, stratified data are collected. The strata may correspond to different habitats, geographic regions or years. If abundance estimates pooled across these categories are required, then the distance data may be pooled across strata for estimating the detection function without generating bias, provided detectability does not vary greatly between strata. This is the pooling robustness property (Buckland *et al.* 2004:389–392). Encounter rate should still be estimated by stratum; the software Distance allows the user to select these options. If separate estimates of abundance are required by stratum, because differences in abundance among strata are of interest, then pooling robustness does not apply. In this case, Distance allows the user to assess whether it is reasonable to assume a common detection function across strata, using AIC. If it is not reasonable, then we can model the detection function as a function of covariates (Marques and Buckland 2003; Marques *et al.* 2007). This allows us to assume a common model across strata, but to allow the parameters determining the scale of the function to vary, depending on recorded covariates. The stratification factor is treated as a factor covariate, with one parameter for each level of the factor.

I do not have enough data to estimate the detection function separately for all the species I surveyed.

This is a frequent problem and is a particular issue when the rarer species with fewer records are of special conservation interest. In this case, we might pool the distance data across groups of species judged to have a similar relationship between detectability and distance so that each species group has sufficient data for estimates of reasonable precision to be made. Marsden *et al.* (1997) used this simple method to gain some indication of the abundance of a rarely recorded endemic fantail *Rhipidura superflua*, by combining its records with the similar and common congeneric *R. rufiventris*. Such abundance figures should be interpreted cautiously because detectability patterns may well differ between the pair of species, and the encounter rate of the rarer species is, by its nature, very low. Allredge *et al.* (2007a) suggest methods for grouping species according to measurements of maximum detection distance and song rate. The approach of the previous section, in which species is specified as a factor covariate and a multiple covariate distance sampling analysis conducted, avoids the assumption that different species have the same

detectability (Marques *et al.* 2007). By treating these factors as main effects, we are assuming that differences in detectability between species are the same in different habitats or years, and that differences in detectability between years are the same for each species. Such assumptions may be known to be reasonable because of information on the sound attenuating properties of different types of vegetation or observer skill. The inclusion of covariates is readily done in Distance and assumptions can be tested, for example by testing for interaction between species and year.

My birds respond to my presence by moving.

It can be difficult to discover that responsive movement is occurring. Attraction to the observer tends to generate too many detections close to the line or point, and is difficult or impossible to distinguish from a more rapid fall-off in detectability with distance. If present, it leads to overestimation of density. Avoidance can sometimes be detected, because there tends to be an excess of detections at mid-distances relative to near the line or point; in the absence of avoidance, this would suggest that detectability increases with distance from the line or point initially, which is implausible for most surveys. Detectability is sometimes lower at the line or point than at a small distance from it, for example in aerial surveys where detection of birds close to the line is difficult, or for ground surveys if birds very close to the point respond to the observer by going silent. More extreme avoidance, in which a proportion of birds move out of the sampled plot before detection, is not evident from the data, and leads to underestimation of density. It may be possible to set up trials with monitored birds, and record their reactions to an approaching observer, to assess the extent of bias from responsive movement.

For some species such as gamebirds, we often rely on responsive movement (flushing) to detect them. This is not a problem, provided their location prior to the flush is known, and the distance from it to the point or line is recorded. Such species should be surveyed using line transect sampling rather than point transect sampling; in the latter case, birds may be flushed when the observer is 'off-effort', walking into the point, and the observer is not moving while recording, so is unlikely to flush birds while 'on-effort'.

If an observer can reach a point without scaring birds away, he or she can wait a while before recording commences, to allow birds to return to their normal activities or original locations. This is called using a settling down period. If this approach is appropriate, point transect sampling may result in less responsive movement than line transect sampling. However, if some species move away from the observer, they may leave the plot during the settling down period and this will lead to underestimation. Some birds become more detectable when disturbed. If they don't move appreciably from their initial location, it may be preferable to start recording as soon as the point is reached. In contrast, if birds temporarily become less detectable but remain in place, then a settling down period should be used. When establishing a new survey, trials might be conducted to assess species' behaviour, from which decisions on settling down periods can be made. Within the same study, some species might be allowed a settling down period and others not. Lee and Marsden (2008) found that a settling down period led to fewer detections of canopy frugivores and upperstorey gleaning insectivores, especially within 10 m of the point, and concluded that corresponding densities were likely to have been underestimated due to evasive movement. In contrast, there were more detections of sallying insectivores when a settling down period was used.

If responsive movement occurs, field procedures should be developed that allow as far as possible detection of birds before they respond to the observer. This may mean concentrating search effort in the forwards direction rather than to the sides in the case of line transect sampling, and the observer should use aids where feasible to allow detection at greater distances; for example in shipboard surveys of seabirds, large tripod-mounted binoculars might be used. For terrestrial line transect surveys, observers walking in line abreast, symmetrically spread either side of the transect line, with a separation distance that ensures detection of any birds that

are located between observers, can reduce responsive movement effects: if birds move away from one observer, they may move towards another; and birds attracted towards observers will sometimes move closer to the transect line and sometimes further from it. This strategy also widens the strip that is effectively surveyed, and generates distance data from which the detection function can be more reliably modelled, because the area where detection is certain is increased.

My birds are detected in two very different ways.

For some bird species such as galliforms, pittas or ground-dwelling pigeons, individuals are detected either visually very close to the transect line or point, or at greater distances when they call sporadically. In such cases, it is often preferable to consider the two types of data separately. For example, for calling birds, cue count methods conducted from lines or points may prove effective. If there are also sufficient visual detections, and the assumptions of distance sampling are believed to hold for these, a line or point transect analysis might be conducted on the sightings data. If both methods yield reliable estimation, the estimates may be combined for example as a weighted average, with weights equal to say the reciprocals of the variances of the estimates.

Is it a problem if I am unable to identify all individuals of the focal species?

Provided that all individuals of the focal species very near the line or point are detected and identified correctly, there is no requirement that more distant birds are all detected or all identified. Records of uncertain identity should be omitted. In this circumstance, the detection function reflects the probability that a bird is both detected and identified, as a function of distance from the line or point.

I have conducted intensive/extensive surveys but still have only a handful of records for a species – how can I quantify this rarity?

If you can also record a species with similar detectability that is more common, you can use multiple covariate distance sampling to analyse the pooled data, including species as a factor covariate, as noted earlier. If this is not an option, but your survey continues for several years, you can again use multiple covariate distance sampling to analyse data pooled across years, with year as either a continuous covariate or a factor covariate in the model for the detection function. Model selection tools such as AIC allow you to assess whether a common detection function can be assumed over time. If yours is a one-off survey, a possible strategy is to locate an individual, then have observers conduct trials to see whether they detect the individual at varying distances. This can be repeated with as many individuals as is feasible, and the resulting estimated detection function (possibly incorporating covariates and fitted for example using logistic regression) might be assumed to apply to the main survey.

How do I test for statistical differences in bird densities across habitats, sites or years?

It is relatively straightforward to test whether two density estimates differ significantly. Provided the design has a reasonable number of lines or points, and provided that the number of detections is adequate to fit a detection function reliably, the degrees of freedom associated with the estimated variances of the density estimates will be sufficient that a large-sample z-test can be used; otherwise, an approximate t-test can be used (Buckland *et al.* 2001:84–85). If the two density estimates have a common estimated detection function, then care is needed to take account of this (Buckland *et al.* 2001:85–86). If there is more complex dependence between the

two estimates, then it is simpler to use nonparametric bootstrap resampling of lines or points to estimate a confidence interval for the difference in true densities. This has the advantage that the estimates do not need to be independent; we eliminate the dependence by estimating the difference in densities within each resample. If zero is included in the bootstrap interval, then the data are consistent with the hypothesis that the two densities are the same.

When there are more than two estimates, and we wish to test whether at least one of them differs significantly from the rest, again a bootstrap approach is perhaps simplest. Under the null hypothesis that all densities are the same, we can estimate this common density by a weighted average of the estimates, where weights are the inverse of the estimated variances. We might then assume that a given estimated density is distributed as log-normal with mean equal to the weighted mean from above, and variance set equal to the estimated variance corresponding to that estimate. A parametric bootstrap resample is generated from this distribution, from which density is estimated. This is repeated for each of the original density estimates. A measure of variability of the density estimates is calculated for the original estimates along with each of the bootstrap sets of estimates. A suitable measure might be the sample variance of the estimates, again weighted by the reciprocal of the estimated variance of each of the original estimates. If the measure calculated from the original estimates is amongst the 5% of most extreme values obtained from the bootstrap resamples, then the null hypothesis that there is a common density is rejected at the 5% level. Thus this is a parametric bootstrap hypothesis test. It can also be regarded as a Monte Carlo test in which the parameters of the assumed distributions are estimated. As with the two-sample bootstrap, the method is unaffected by any dependence amongst the estimated densities, such as arises if there is a common model for the detection function.

References

- Allredge, M. W., Pollock, K. H., Simons, T. R. and Shriener, S. A. (2007a) Multiple-species analysis of point count data: a more parsimonious modelling framework. *J. Appl. Ecol.* 44: 281–290.
- Allredge, M. W., Simons, T. R. and Pollock, K. H. (2007b) A field evaluation of distance measurement error in auditory avian point count surveys. *J. Wildlife Manage.* 71: 2759–2766.
- Bibby, C. J. and Buckland, S. T. (1987) Bias of bird census results due to detectability varying with habitat. *Acta Oecologica: Oecol. Gener.* 8: 103–112.
- Bibby, C. J., Burgess, N. D., Hill, D. A. and Mustoe, S. H. (2000) *Bird census techniques*. 2nd edn. London: Academic Press.
- Borchers, D. L., Laake, J. L., Southwell, C. and Paxton, C. G. M. (2006) Accommodating unmodeled heterogeneity in double-observer distance sampling surveys. *Biometrics* 62: 372–378.
- Brandes, T. S. (2008) Automated sound recording and analysis techniques for bird surveys and conservation. *Bird Conserv. Internatn.* 18: S163–S173.
- Buckland, S. T. (1987) On the variable circular plot method of estimating animal density. *Biometrics* 43: 363–384.
- Buckland, S. T. (2006) Point transect surveys for songbirds: robust methodologies. *Auk* 123: 345–357.
- Buckland, S. T., Anderson, D. R., Burnham, K. P., Laake, J. L., Borchers, D. L. and Thomas, L. (2001) *Introduction to distance sampling*. Oxford: Oxford University Press.
- Buckland, S. T., Anderson, D. R., Burnham, K. P., Laake, J. L., Borchers, D. L. and Thomas, L., eds. (2004) *Advanced distance sampling*. Oxford: Oxford University Press.
- Buckland, S. T., Summers, R. W., Borchers, D. L. and Thomas, L. (2006) Point transect sampling with traps or lures. *J. App. Ecol.* 43: 377–384.
- Cabanne, G. S., Zurita, G. A., Seipke, S. H. and Bellocq, M. I. (2007) Range expansion, density and conservation of the Araucaria Tit-spinetail *Leptasthenura setaria* (Furnariidae) in Argentina: the role of araucaria *Araucaria angustifolia*

- (Araucariaceae) plantations. *Bird Conserv. Internatn.* 17: 341–349.
- Conti Nunes, M. F., Galetti, M., Marsden, S., Pereira, R. S. and Peterson, A. T. (2007) Are large-scale distributional shifts of the blue-winged macaw (*Primolius maracana*) related to climate change? *J. Biogeog.* 34: 816–827.
- Farnsworth, G. L., Pollock, K. H., Nichols, J. D., Simons, T. R., Hines, J. E. and Sauer, J. R. (2002) A removal model for estimating detection probabilities from point-count surveys. *Auk* 119: 414–425.
- Freeman, S. N., Noble, D. G., Newson, S. E. and Baillie, S. R. (2007) Modelling population changes using data from different surveys: the Common Birds Census and the Breeding Bird Survey. *Bird Study* 54: 61–72.
- Githiru, M. and Lens, L. (2006) Demography of an Afrotropical passerine in a highly fragmented landscape. *Anim. Conserv.* 9: 21–27.
- Greenwood, J. J. D. and Robinson, R. A. (2006) Principles of sampling. Pp. 11–86 in W. J. Sutherland, ed. *Ecological census techniques*. 2nd edn. Cambridge: Cambridge University Press.
- Gregory, R. D., Baillie, S. R. and Bashford, R. I. (2000) Monitoring breeding birds in the United Kingdom. *Bird Census News* 13: 101–112.
- Hedley, S. L., Buckland, S. T. and Borchers, D. L. (2004) Spatial distance sampling models. Pp. 48–70 in S. T. Buckland, D. R. Anderson, K. P. Burnham, J. L. Laake, D. L. Borchers and L. Thomas, eds. *Advanced distance sampling*. Oxford: Oxford University Press.
- Hiby, A. R. (1985) An approach to estimating population densities of great whales from sighting surveys. *IMA J. Math. Appl. Med.* 2: 201–220.
- Hilton, G. M., Atkinson, P. W., Gray, G. A. L., Arendt, W. J. and Gibbons, D. W. (2003) Rapid decline of the volcanically threatened Montserrat oriole. *Biol. Conserv.* 11: 79–89.
- IUCN (2001) *IUCN Red List Categories and Criteria: Version 3.1*. Gland, Switzerland and Cambridge, UK: IUCN Species Survival Commission.
- Jepson, P., Brickle, N. and Chayadin, Y. (2001) The conservation status of Tanimbar corella and blue-streaked lorry on the Tanimbar Islands, Indonesia: results of a rapid contextual survey. *Oryx* 35: 224–233.
- Jones, M. J., Marsden, S. J. and Linsley, M. D. (2003) Effects of habitat change and geographical variation on the bird communities of two Indonesian islands. *Biodivers. Conserv.* 12: 1013–1032.
- Laake, J. L. and Borchers, D. L. (2004) Methods for incomplete detection at distance zero. Pp. 108–189 in S. T. Buckland, D. R. Anderson, K. P. Burnham, J. L. Laake, D. L. Borchers and L. Thomas, eds. *Advanced distance sampling*. Oxford: Oxford University Press.
- Lambert, F. R. (1993) Trade, status and management of three parrots in the north Moluccas, Indonesia: White Cockatoo *Cacatua alba*, Chattering Lory *Lorius garrulus* and Violet-eared Lory *Eos squamata*. *Bird Conserv. Internatn.* 3: 145–168.
- Lee, D. C. and Marsden, S. J. (2008) Adjusting count period strategies to improve the accuracy of forest bird abundance estimates from point transect distance sampling surveys. *Ibis* 150: 315–325.
- Maldonado-Coelho, M. and Marini, M. A. (2004) Mixed-species bird flocks from Brazilian Atlantic forest: the effects of forest fragmentation and seasonality on their size, richness and stability. *Biol. Conserv.* 116: 19–26.
- Marques, F. F. C. and Buckland, S. T. (2003) Incorporating covariates into standard line transect analyses. *Biometrics* 59: 924–935.
- Marques, T. A. (2007) *Incorporating measurement error and density gradients in distance sampling surveys*. PhD thesis, University of St Andrews, UK.
- Marques, T. A., Thomas, L., Fancy, S. G. and Buckland, S. T. (2007) Improving estimates of bird density using multiple covariate distance sampling. *Auk* 124: 1229–1243.
- Marsden, S. J. (1999) Estimation of parrot and hornbill densities using a point count distance sampling method. *Ibis* 141: 377–390.
- Marsden, S. J., Jones, M. J., Linsley, M. D., Mead, C. and Hounscome, M. V. (1997) The conservation status of the restricted-range

- lowland birds of Buru, Indonesia. *Bird Conserv. Internatn.* 7: 213–233.
- Marsden, S. J., Whiffin, M., Galetti, M. and Fielding, A. H. (2005) How well will Brazil's network of lowland Atlantic forest reserves maintain viable bird populations? *Biodivers. Conserv.* 14: 2835–2853.
- Marsden, S. J., Symes, C. T. and Mack, A. L. (2006) The response of a New Guinean avifauna to conversion of forest to small-scale agriculture. *Ibis* 148: 629–640.
- Mulwa, R. K., Bennun, L. A., Ogot, C. K. P. O. and Lens, L. (2007) Population status and distribution of Taita White-eye *Zosterops silvanus* in the fragmented forests of Taita Hills and Mount Kasigau, Kenya. *Bird Conserv. Internatn.* 17: 141–150.
- Newson, S. E., Evans, K. L., Noble, D. G., Greenwood, J. J. D. and Gaston, K. J. (in press) Estimates of national population sizes for common and widespread breeding birds in the UK. *J. Appl. Ecol.*
- Newson, S. E., Woodburn, R. J. W., Noble, D. G., Baillie, S. R. and Gregory, R. D. (2005) Evaluating the Breeding Bird Survey for producing national population size and density estimates. *Bird Study* 52: 42–54.
- Nichols, J. D., Hines, J. E., Sauer, J. R., Fallon, F. W., Fallon, J. E. and Heglund, P. J. (2000) A double-observer approach for estimating detection probability and abundance from point counts. *Auk* 117: 393–408.
- Staicer, C. A., Ingalls, V. and Sherry, T. W. (2006) Singing behavior varies with breeding status of American Redstarts (*Setophaga ruticilla*). *Wilson J. Ornithol.* 118: 439–451.
- Strindberg, S., Buckland, S. T. and Thomas, L. (2004) Design of distance sampling surveys and Geographic Information Systems. Pp. 190–228 in S. T. Buckland, D. R. Anderson, K. P. Burnham, J. L. Laake, D. L. Borchers and L. Thomas, eds. *Advanced distance sampling*. Oxford: Oxford University Press.
- Thomas, L., Laake, J. L., Strindberg, S., Marques, F. F. C., Buckland, S. T., Borchers, D. L., Anderson, D. R., Burnham, K. P., Hedley, S. L., Pollard, J. H., Bishop, J. R. B. and Marques, T. A. (2006) *Distance 5.0*. University of St. Andrews, UK: Research Unit for Wildlife Population Assessment. <http://www.ruwpa.st-and.ac.uk/distance/>
- van Dongen, W. F. D. (2006) Variation in singing behaviour reveals possible functions of song in male golden whistlers. *Behaviour* 143: 57–82.

STEPHEN T. BUCKLAND*

Centre for Research into Ecological and Environmental Modelling, University of St Andrews, The Observatory, Buchanan Gardens, St Andrews KY16 9LZ, U.K.

STUART J. MARSDEN

Applied Ecology Group, Dept of Environmental and Geographical Sciences, Manchester Metropolitan University, Chester Street, Manchester M1 5GD, U.K.

RHYS E. GREEN

Royal Society for the Protection of Birds and Conservation Ecology Group, Dept of Zoology, University of Cambridge, Downing Street, Cambridge CB2 3EJ, U.K.

*Author for correspondence; e-mail: steve@mcs.st-and.ac.uk