On the Analysis of a Cloud Seeding Dataset over Tasmania

ANTHONY E. MORRISON, STEVEN T. SIEMS, AND MICHAEL J. MANTON

Monash University, Melbourne, Victoria, Australia

ALEX NAZAROV

Hydro Tasmania, Ltd., Hobart, Tasmania, Australia

(Manuscript received 26 June 2008, in final form 30 October 2008)

ABSTRACT

An analysis of cloud seeding activity for the period 1960–2005 over a hydroelectric catchment (target) area located in central Tasmania, Australia, is presented. The analysis is performed using a double ratio on monthly area-averaged rainfall for the months of May–October. Results indicate that increases in monthly precipitation are observed within the target area relative to nearby controls during periods of cloud seeding activity. Ten independent tests were performed and all double ratios found are above unity with values that range from 5% to 14%. Nine out of 10 confidence intervals are entirely above unity and overlap in the range of 6%–11%. Nine tests obtain levels of significance >0.05 level. If the Bonferroni adjustment is made to account for multiple comparisons, six tests are found to be significant at the adjusted alpha level. Further field measurements of the cloud microphysics over this region are needed to provide a physical basis for these statistical results.

1. Introduction

The practice of cloud seeding has remained a point of contention in the scientific community for over half of a century. Early laboratory experiments were able to readily demonstrate precipitation enhancement mechanisms through the conversion of supercooled water to ice by the introduction of suitable ice nuclei (Schaefer 1946), and these laboratory experiments were followed by a field demonstration on individual clouds by Kraus and Squires (1947). However, the extension of cloud seeding impacts from individual clouds to a sustained precipitation increase over a substantial surface area has proven to be an elusive goal, especially at the high level of proof required by the wider scientific community.

The 2003 U.S. National Research Council (NRC) report, titled ''Critical issues in weather modification research,'' includes an abridged history of the development of various methods of cloud seeding with numerous references to static glaciogenic (of both cumulus and winter orographic regimes), dynamic glaciogenic, and hygroscopic seeding field experiments. Despite these considerable research efforts spanning decades, the report goes on to highlight the persistence of key uncertainties, which are broadly classified as ''cloud/precipitation microphysics issues, cloud dynamic issues, cloud modeling issues and seeding related issues.'' The NRC report ultimately concludes that ''there still is no convincing scientific proof of the efficacy of intentional weather modification efforts.''

Boe et al. (2004) noted that the definition of ''convincing scientific proof'' was ambiguous, leading to Garstang et al. (2005) further clarifying that scientific proof was defined as an understanding of ''processes that can be replicated by predictable, detectable and verifiable results.'' Ultimately, on the question of verification it was recognized that ''the level of noise in natural systems compared to the magnitude of the signal has made verification of either the enhancement of rain or snowfall or the reduction of hail extremely difficult.''

In principle it is possible to overcome large variability by extending a trial so that the accepted 5% significance level can be achieved. In practice it is not clear what would constitute a suitable period given that precipitation shows variability on the time scale of hours to decades.

Corresponding author address: Anthony Morrison, School of Mathematical Sciences, Monash University, Wellington Rd., Clayton, Melbourne, VIC 3800, Australia.
E-mail: anthony.morrison@sci.monash.edu.au

DOI: 10.1175/2008JAMC2068.1

A very practical time limit arises over the ability to maintain a consistent, extended scientific experiment. Finite funding, changing personnel, changing technology, and a changing environment all serve to prohibit field work from being sustained over decades. Long-running cloud seeding projects, such as in Israel (Gabriel and Rosenfeld 1990; Nirel and Rosenfeld 1995), the Sierra Nevada (Reynolds and Dennis 1986; Deshler et al. 1990), and Thailand (Silverman and Sukarnjanasat 2000), can become operational, making it difficult to observe and quantify any positive effect over extended periods of time. Moreover, statistical significance is not sufficient to provide ‘‘acceptable proof’’: associated physical observations of expected changes in cloud properties need to be documented to complement any statistical evaluation.

Australia, the driest inhabitable continent, invested heavily in cloud seeding research in the 1980s. As in the United States, funding for such research all but ceased at this time because of a lack of convincing scientific proof. Ryan and King (1997) present an account of the many cloud seeding research programs dating back to 1947 conducted by the Commonwealth Scientific and Industrial Research Organization (CSIRO). A remarkable result reported in the literature was over the island of Tasmania (Smith et al. 1979), where cloud seeding was found to produce a statistically significant increase in surface rainfall over a target of approximately 2500 km^2. Specifically, a 30% increase was found in the autumn months (April, May, and June) with weaker increases (that failed to reach high levels of significance) observed in the winter months and early spring. Given the inability of CSIRO to produce such positive results elsewhere [except in the Snowy Mountains (Smith et al. 1963)], the results were met with some widespread skepticism, leading to a second Tasmanian cloud seeding experiment being conducted a decade later. For this second field experiment Ryan and King (1997) report a 37% increase in surface rainfall over the months of April–October under specific synoptic conditions. Section 2 presents a thorough review of both of these field experiments and the ensuing operational periods.

Despite two positive field experiments (hereinafter Tasmania 1 and 2), cloud seeding in Tasmania still fails to meet the standard of convincing scientific proof among the wider scientific community for a variety of reasons. Many of the key uncertainties identified in the NRC report are well illustrated in these experiments. There was no ability to directly observe the in situ microphysical response to the cloud seeding as was done in Deshler et al. (1990) within orographic clouds over the Sierra Nevada or in Rosenfeld and Woodley (1989) in summer convective clouds over west Texas; for example, remote sensing technology such as radar or satellite observations was unavailable. The definition of controls employed in Tasmania 1 and the changes to the controls for Tasmania 2 opened legitimate questions about the robustness of the statistical formulation. Indeed, simply the poor documentation of the synoptic meteorology has led to some confusion: the 2003 NRC report lists these Tasmanian experiments as examples of winter orographic cloud seeding, while a closer examination of the meteorology indicates that the situation in Tasmania fails to meet this classification.

These key questions about Tasmania 1 and 2 cannot be addressed some 25 years after the latest fieldwork was undertaken. However, the extended but intermittent record of cloud seeding over the central plateau of Tasmania creates a unique dataset that has the potential to offer insight into the effectiveness of long-term cloud seeding in this region. The notion is further encouraged by the fact that over the last 50 years the catchment area has remained relatively untouched with respect to human influences and economic development; it is protected as the Tasmanian Wilderness World Heritage Area.

From the period of 1960–2005, some form of cloud seeding has taken place during 24 of the 46 winter seasons. If the reported 30% increase found in the first Tasmanian experiment is accepted at face value then it is reasonable to expect a positive signal in the monthly rain gauge records. For example, if 90% of the rainfall was to occur in the 10 wettest days of a given month and half of these days were seeded, then one would expect to see a 13.5% increase in the overall monthly rainfall. Simply seeding any five random days within a month with a 30% increase should still lead to an increase of 5% in the overall monthly rainfall. If seeding were largely ineffective, then one might expect to find no positive effect in the monthly records. With a long enough time record, any meaningful signal should be detectable and found to be statistically significant.

The objectives of this paper are first to analyze the monthly rainfall records over Tasmania for a detectable signal with respect to the act of cloud seeding, and second to evaluate whether any such signal is statistically significant. As the various seeding periods were never designed to be analyzed as a single, long-running time series, numerous caveats exist to this approach. In section 2 a review of the various seeding efforts over central Tasmania is presented. This is followed by a discussion of the meteorology and climatology in section 3. In section 4 the preparation of the rainfall observations is detailed. In section 5 the double ratio technique is employed to quantify a signal with respect to cloud seeding. A bootstrap analysis is then undertaken

TABLE 1. Seeding history.

Seeding period	Mode	Seeding agent	No. of winters seeded
1964–71	Research	Silver iodide	5
1979–83	Research	Silver iodide	5
1988–91	Operational	Silver iodide	4
1992–94	Research	Dry ice	3
1998–2005	Operational	Silver iodide	7

to establish the statistical significance of this signal. Results are more fully discussed in section 6.

2. Historical review of cloud seeding in Tasmania

The island of Tasmania has been the target of both experimental and operational cold cloud seeding dating back to the 1960s (Table 1). Glaciogenic seeding research projects led by CSIRO were conducted between 1964 and 1971 and between 1979 and 1983 [Tasmania 1 and 2 as defined in Ryan and King (1997)]. A third trial was conducted between 1992 and 1994 solely by the island's hydroelectric energy company Hydro Tasmania (HT), formally the Hydro Electric Commission of Tasmania (HEC), although no results from this research period have been published. Two periods of operational cloud seeding have also been conducted, from 1988 to 1991 and 1998 to present.

a. Tasmania 1: 1964, 1966, 1968, 1970–71

Tasmania 1 was conducted as a randomized experiment during all months of the year, over a target area of approximately 2500 km^2 located in central Tasmania. This original target area is referred to as the CSIRO target. The primary analysis was defined as the double ratio of target rainfall relative to a number of control areas: northwest, north, and southeast of the target area (Smith et al. 1979; see Fig. 1 for a map showing the location of these various regions). Seeding units were defined in pairs of duration of 12–18 days, with each pair having the seeded and nonseeded part randomly assigned. The analysis was separated into seasons and implemented independently on both a western and eastern half of the target area. The seeding agent used at that time was an acetone solution of silver and sodium iodide and was released by a single aircraft 30 min upwind (as defined at the seeding level) of the target area for cumulus clouds and 45 min upwind for stratiform clouds. The criteria used in defining suitable conditions were that cloud tops contain supercooled liquid water (SLW) at a temperature colder than −5°C for stratiform clouds and −10°C for cumuliform clouds. Furthermore, clouds had to be deep, compact, and without excessive clear-air volumes. At this time, reliable quantitative instruments for measuring the liquid content of a cloud were not available, but it was noted in Smith et al. (1979) that airframe icing was usual and often severe, occurring about three-quarters of the seeding time. The average seeding time was close to 8 h per month for the whole experiment; however, it is noted that this figure increased steadily from 3.2 h for months during the initial year to 14.2 h during the final year.

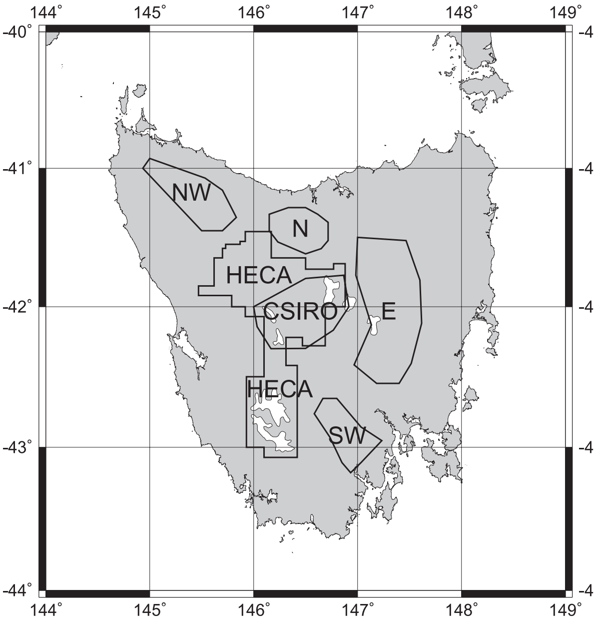

FIG. 1. Map showing the locations of the CSIRO and HECA target areas. Also shown are the static controls used in Tasmania 1 and 2.

Precipitation increases of up to 30% during autumn (March, April, and May) were published in Smith et al. (1979) as defined by the double ratio between the eastern half of a target area and an average of two controls based to the north and southeast of the target area. Weaker evidence was present indicating a probable 23% increase in the western half of the target area during autumn and a possible 13% increase in winter; however, these results were never published in the reviewed literature as they did not reach the required level of significance. Interestingly, it was noted that the observed increase occurred when deep prefrontal stratiform clouds were present, rather than orographically forced clouds, which are commonly studied with respect to static glaciogenic cloud seeding (Long and Huggins 1992; Deshler et al. 1990; Rauber and Grant 1987).

b. Tasmania 2: 1979–83

Based on knowledge gained during Tasmania 1, the operations and science plan for Tasmania 2 were revised in a number of respects. Seeding was limited strictly to the months of April–September with a randomization scheme having a seed/no-seed ratio of 2:1. The seeding unit was reduced from 12–18-day blocks to a single calendar day. Further constraints were added to the definition of a suitable day (i.e., for stratiform clouds, cloud-top temperature $\leq -5°C$, depth is 1/3 terrain clearance of base, supercooled water at the seeding level >0.1 g m^{-3} or, failing that, 2 mm of ice accreted in 5 min on an icing rod of 2-mm diameter and winds at seeding height <130 km h^{-1}). For cumulus clouds, the constraints are cloud top colder than $-12°C$, depth greater than height of base, supercooled water at the $-10°C$ level >0.5 g m^{-3} or alternatively 1 mm of ice accreted on an icing rod 2 mm in diameter during one pass, bases flat and ''firm'' with tops extending vertically, and wind speed at cloud base <100 km h^{-1}. It is noted in Shaw et al. (1984) that this tightening up of criteria drastically reduced the number of suitable days available for seeding. For example, the average number of suitable days during autumn and winter was 52 for Tasmania 1. In Tasmania 2, the average was 18. Further modifications to the experimental procedure were that stratiform clouds should be seeded 1 h upwind instead of 45 min as was the case in Tasmania 1 and that the seeding solution be modified from silver and sodium iodide to silver and ammonium iodide because of improved nucleating abilities at warmer temperatures.

By 1979 it was possible to quantify cloud water content, so Tasmania 2 was able to record the microphysical mixed phase conditions that were actually encountered by the aircraft. Shaw et al. (1984) and Ryan and King (1997) state that the most frequent liquid water content measured with a 5-min time constant was $\sim$0.3g m^{-3}, occurring between $-6°$ and $-8°C$. This is a large amount of supercooled water relative to studies like Deshler and Reynolds (1990), where the most common peak supercooled water content (per flight track) was $\sim$0.1 g m^{-3} in a similar temperature range.

While the target area defined did not change between Tasmania 1 and 2, a second set of floating controls was introduced for the evaluation of Tasmania 2 in addition to the original fixed controls used in Tasmania 1. The motive for the introduction of floating controls was based on previous Australian cloud seeding experiments where it was found that correlations between rainfalls depended on wind direction. On any experimental day, three floating controls were defined: 1) directly upwind of the target, 2) left of the target when looking into the wind, and 3) right of the target. The final dataset consisted of 66 days for which the average seeding time per month was $\sim$2.7 h. The results were that there was good evidence indicating an effect in the western half of the target area and some evidence indicating a possible effect in the eastern half. If the full 66 days are included in the analysis using the double ratio to estimate the change in rainfall due to seeding, a 30% increase is observed with a P value of 0.007 in the west target. Relative to the fixed controls the estimated increase is only 12% and is not significant. If the analysis is restricted to days where the wind direction is between 231° and 300° (the preferred sector), the estimated increase relative to the floating controls increases to 37%.

c. *Recent cloud seeding*

Given the two positive field experiments reported by CSIRO, the HEC undertook operational cloud seeding in the years 1988–91. The same seeding guidelines were employed for this period as with Tasmania 2, except that no randomization was undertaken and no efforts to quantify results were made. The decision to seed was made in response to low reservoir levels in the hydroelectric catchment area.

The third trial (1992–94), according to Ryan and King (1997), was run solely by the HEC in a similar manner to that of Tasmania 2. Results for this trial have not been published; however, there were a total of 61 seeded flights over the 3-yr period with an average of 5 hours of seeding per month from May–November (log files from HT). The seeding agent used during this period was dry ice.

The current operational program has been running since September 1998. In essence it has many similarities to Tasmania 2 in that clouds are usually only seeded if they meet criteria for supercooled water, temperature, wind speed, and concentrations of ice crystals. However, these are less strict than those in Tasmania 2. The decision to fly is based on satellite observations of cloud-top temperature. Operations are often undertaken in postfrontal conditions to avoid flying when lightning is present, which has been found to be more common in prefrontal systems. A significant difference between the recent years of operational seeding and earlier research activities is that the target area has been expanded from the original CSIRO target to the new hydroelectric catchment area (HECA; Fig. 1). This is an increase of around 5000 km^{2}. Because of the increased size of the region intended for precipitation augmentation, individual seeding operations now focus on specific catchments within the HECA. A typical seeding flight would target up to two of these individual catchment areas.

Suitable clouds are predominantly associated with cold fronts moving across the Southern Ocean and are usually tracked continuously from April to November. Numerical products are supplied by CSIRO in the form of Conformal-Cubic Global Atmospheric Model (CCAM) output (McGregor and Dix 2007) and include height at $-10°C$, total cloud cover, average cloud water and ice, and magnitude of vertical wind shear. Cloud microphysical properties are assessed in flight using a CSIRO King probe for liquid water and a Droplet Measurement Technologies (DMT) Cloud Aerosol and Precipitation Spectrometer (CAPS) probe (Baumgardner et al. 2001) for measuring particle numbers. The decision to seed is based on there being supercooled liquid water present for extended periods of time and minimal ice. Only about one out of four flights actually meets this last criterion (and is seeded). On average there are 32 seeded flights per year between April and November with an average of 5 h of seeding per month.

3. Climate and meteorology

The meteorological conditions under which cloud seeding is performed in Tasmania are determined by the passage of fronts in the Southern Ocean. Tasmania is situated between 40° and 44°S and year-round receives a large fraction of its precipitation from frontal cloud systems as they sweep east over southern Australia. Mace et al. (2007) determined that the Southern Ocean's (which covers nearly 15% of the earth's surface) cloud climatology is unlike any other part of the globe with the largest fraction of low-level cloud cover and a fair portion of relatively thick midlevel cloud. The fractional hydrometeor cover is between 0.7 and 0.8 in an almost continuous cloud band that circles the entire planet. Between 50 and 70 cold fronts usually pass over Tasmania every year suspended in an extremely clean air mass. Cloud condensation nuclei (CCN) concentrations measured at Cape Grim on the northwest coast of Tasmania indicate around 70 CCN cm^{-3} active at 0.23% supersaturation (Gras 1995). SLW is frequently observed in these clouds in both the pre- and postfrontal air mass (Long and Huggins 1992; Ryan and King 1997).

Tasmania can broadly be described as having a fairly mountainous west coast with peaks that reach approximately 1.5 km, a central plateau (where the CSIRO target is mostly located) that has an elevation of close to 1 km, and a relatively low lying eastern region. The extent of the island is around 250 km in the north–south direction and 300 km at the widest point east–west. As a front passes over Tasmania from west to east, precipitation often occurs over the western and central regions. The year-round monthly average precipitation on the west coast is 180 mm, with a low of 100 mm in February and high of 250 mm during July. The central plateau receives between 90 and 180 mm per month and the eastern region between 50 and 80 mm.

Figure 2 shows an example of satellite retrieved cloud-top temperature (CTT) from the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument on the *Aqua* and *Terra* satellites for a front to the west of Tasmania (Fig. 2a) at 0455 UTC 8 August 2006 and to the east (Fig. 2b) at 2350 UTC 8 August 2006. A cloud seeding mission departed Hobart at 2220 UTC and seeded the clouds over central western Tasmania shown in Fig. 2b. This frontal cloud band had been tracked for days in advance and the decision to fly was based on lightning being absent in the postfrontal cloud mass and there being clouds with temperatures between 0° and $-10°C$. The decision to seed when flying was based on the fact that approximately 20% of the in cloud conditions had SLW contents >0.5 g m^{-3} at the seeding altitude. This seeding mission was considered to have fair conditions; it is not unusual to have SLW contents that are higher than this.

Figure 3 shows the locations of the seeding track relative to the HECA for the period 2002–05. It is observed that over the majority of time during seeding conditions, the prevailing wind direction is from either the west, northwest, or southwest; however, approximately 25% of the time the wind direction is not in these sectors. The northwest conditions are typically prefrontal.

4. Data preparation

The purpose of any useful cloud seeding project is to increase annual or at least seasonal rainfall to a measurable extent over a region of economic significance. This investigation attempts to identify a significant increase in rainfall due to cloud seeding by taking a calendar month as the basic unit of time. A monthly averaged rainfall is defined as the arithmetic mean from all long-lived surface rain gauge sites operating anywhere within a specified region. If a station has any missing records within the given calendar month, that specific station is discarded for the entire month.

A shortcoming of using this minimalist definition of an area-averaged rainfall is that the spatial distribution of rain gauges may not be adequate to properly sample the spatial variability of rainfall in the target area, and so the estimated rainfall may not be a fair representation of the area average. However, a comparison of the site-averaged values with a grid-based estimate from the Commonwealth Bureau of Meteorology provides some confidence in this simple approach.

Figure 4 shows that a number of rain gauges have only become operational since 1998, particularly those maintained by HT (for completeness, precipitation totals are also shown). These new HT records contain no information on nonseeded periods and thus offer no

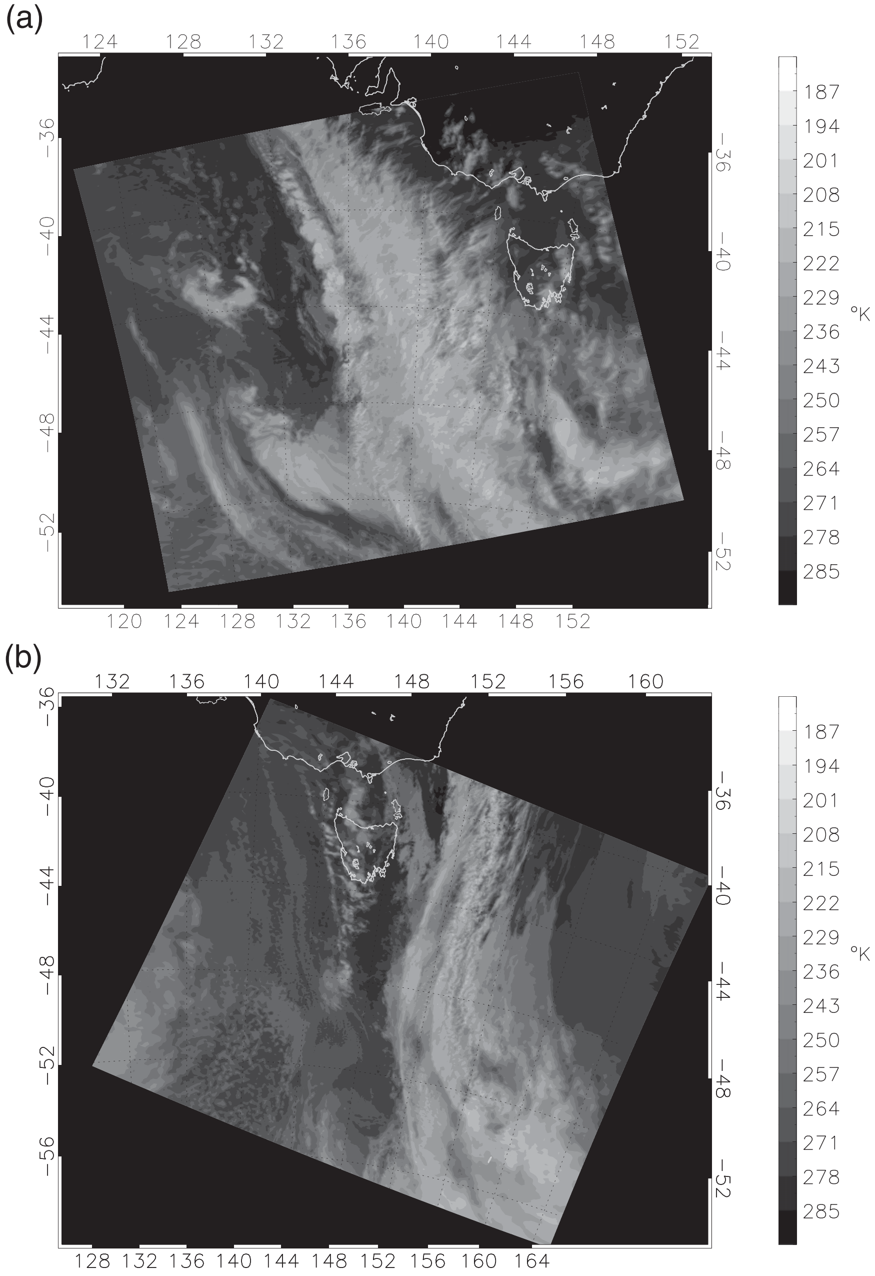

FIG. 2. Satellite retrieved CTT (K). (a) The front to the west of Tasmania and (b) the front to the east of Tasmania.

insight into the science. As a means of removing these (or other) short-lived sites, a longevity filter is used on individual sites. Each individual site was filtered at three levels: those having operated for a minimum of 25%, 50%, and 75% of the total time from 1960 to 2005. This filtering is independent of the act of seeding. As the various tests were found to be relatively insensitive to this filter, all discussion from here on will refer to the 50% threshold data unless otherwise specified. This longevity filter was chosen because it represented a balance between adequate spatial sampling with regard to the number of gauges and minimal temporal evolution of the rain gauge network (i.e., the coming and going of individual sites). To further expand this, the analysis was also performed using only the Bureau of Meteorology sites, as there may be a perception that these sites are of a higher standard. The omission of the HT sites has no qualitative impact or systematic bias.

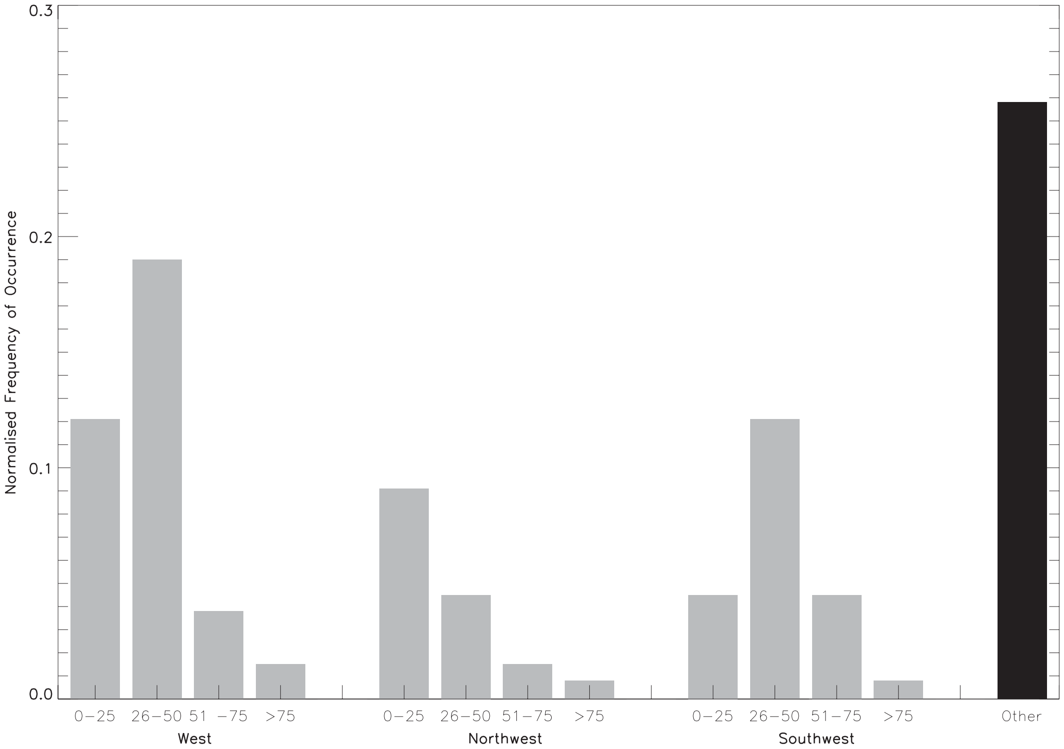

FIG. 3. Frequency of occurrence indicating the approximate position of the seeding track relative to the target area: west, northwest, southwest, or other for the period 2002–05. Four distances are shown: 0–25, 26–50, 51–75, and >75 km. The column labeled other contains no data (i.e., not documented) or tracks anywhere else over Tasmania.

Using a calendar month as a seeding unit has a number of advantages. First, there are no ambiguities in the definition of seeded periods; if the aircraft made even a single seeding mission in a given month, that calendar month is a seeded month. Second, the use of a broad temporal average greatly reduces the variability of the dataset, which should allow for statistical significance to be more easily established.

It is noted that months with only one or two seeding flights are retained as seeded months, even though it is difficult to imagine that such limited seeding could have any impact on monthly rainfall totals. Indeed, these lightly seeded months have a lower average rainfall than the overall average for the 46-yr period. This is not surprising, as these lightly seeded months are likely to have had few seeding opportunities and thus correspondingly poor rainfall. They must be retained as seeded months, however, to prevent a bias of the data toward months with wetter conditions.

A nonseeded month primarily means that the aircraft was not in operation, regardless of the meteorology. The seeding history reveals that even during periods of poor regional rainfall, the aircraft will still have seeded at least once in a given month if it were available. For example, there is not a single nonseeded winter month in the most recent operational period, even though it has been widely reported that southeast Australia has been suffering from a drought over this period (although this is not the case for some regions in Tasmania) (Watkins and Trewin 2007; National Climate Centre 2006). Over the full seeding history, there are only seven nonseeded winter months (June–August) immediately adjacent to a seeded month. It is likely that these missed months were primarily a result of operational rather than meteorological factors. In contrast, there are 16 winter months that have been seeded only a single time, and 10 winter months that have been seeded only twice.

Only months with more than 60 flights are included in the analysis, so that the treatment of months is consistent. Figure 5 shows the number of seeded months and seeded missions for 1960–2005. It is observed that almost no seeding occurs in the summer and early autumn (December–March). As discussed in section 2, seeding was only attempted in the summer period during Tasmania 1. These summer months are excluded in the present analysis, simply because of undersampling with respect to the seeded sample. While the months of April and November have a greater number of seeded units (relative to the summer months), these are also excluded

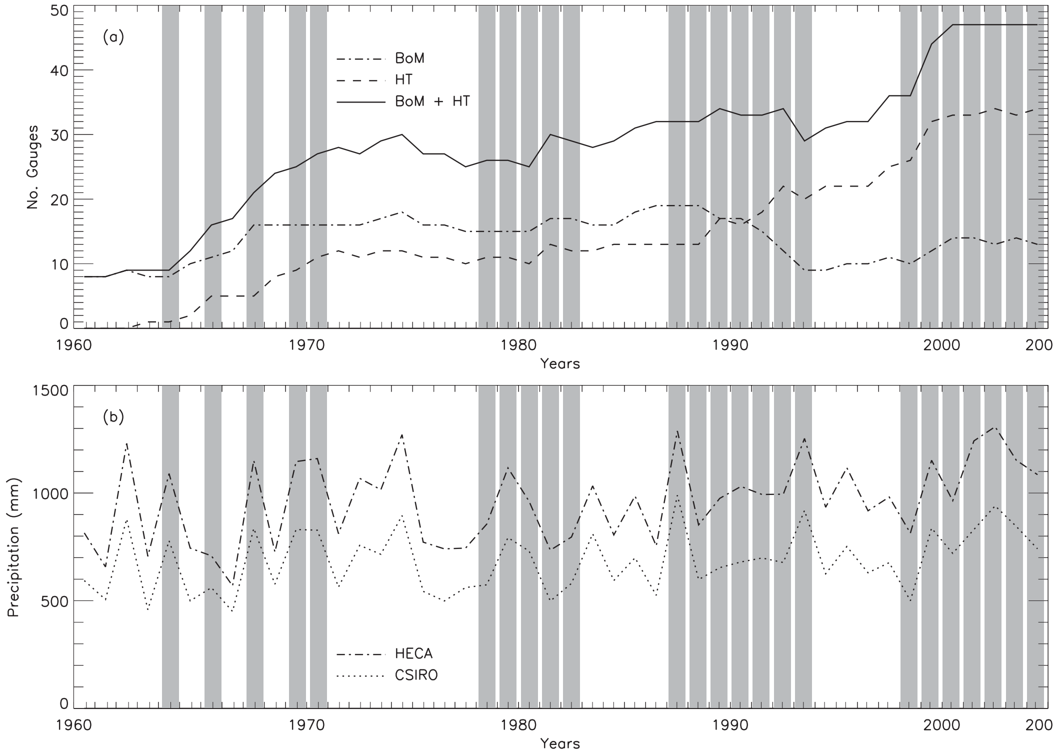

FIG. 4. (a) Time series of the number of rain gauges operating in the HECA target area from 1960 to 2005. (b) Quantity of precipitation that fell during May–October 1960–2005 for both the HECA and CSIRO targets. Years in which winter seeding occurred are shaded. Note that the average number of sites used to construct the HECA (CSIRO) area average is 20 (12).

from the present analysis. During operational periods if the aircraft was ready to fly during late April and suitable conditions were present, seeding missions would have begun early. Similarly, November is omitted as yearly operations were often ended if meteorological conditions were poor early in the month. Hence, the inclusion of these months could introduce a bias toward favorable meteorological conditions. Note that October has fewer seeding months than May–September, but this is only because October was explicitly not seeded during Tasmania 2 (i.e., the decision was operational rather than meteorological).

Limiting the analysis to the months May–October 1960–2005, the dataset now encompasses a total of 276 months of which 130 have been seeded. Table 2 shows the average rainfalls for both the CSIRO target and the HECA. As detailed in section 2, the target area has incrementally expanded from an original size of 2500 km^2 (CSIRO) to 7500 km^2 (HECA) over the periods of operational cloud seeding. Qualitatively the average rainfall in the two targets behaves similarly over the duration of the experiment (Fig. 4). As the western portion of the HECA target is at a higher elevation than the original CSIRO target, rain gauges in this area experience a greater rainfall on average. Thus while it is possible to define a time-dependent, expanding target region, such sophistication raises ambiguity. A more straightforward approach is to simply analyze the two limits for the target region.

The difference in elevation (and observed rainfall) across the target combined with the time-dependent nature of the rain gauge network could introduce a further bias. For example, a rain gauge operating at high elevation during limited periods would lead to greater area-average rainfalls during those periods, whether seeded or not. While the longevity filter already in place should act to minimize this type of bias, this effect must still be explicitly assessed. The test designed to assess this potential caveat involved substituting the observed monthly rainfall for each rain gauge with idealized data unbiased with regard to the act of seeding.

Specifically, each individual rain gauge is interrogated for its rainfall totals for all months having a complete set of data. These data were then compared with a number of ideal distributions. It was found that a lognormal distribution most accurately represented the real data with approximately 80% of the months exceeding the 0.05 significance level, tested using the

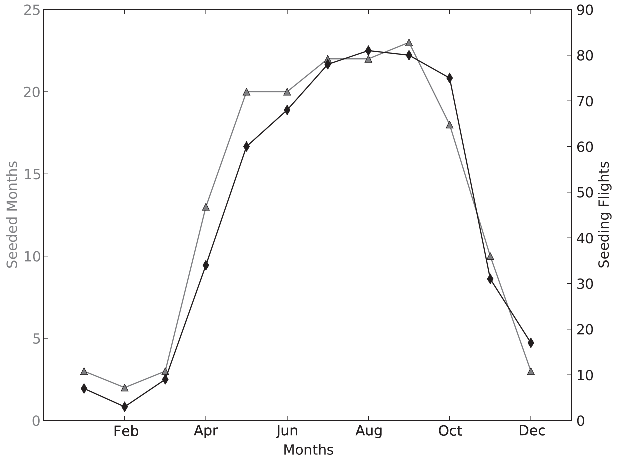

FIG. 5. Number of seeding events per month (diamonds) and number of months seeded (1960–2005) (triangles).

Jarque–Bera test of the null hypothesis that a given sample comes from a specific distribution (Jarque and Bera 1980). Idealized random data were then created for the monthly rainfall for each site. These data were then analyzed for the same values as seen in Table 2. One thousand datasets were created and analyzed for the period of 1960–2005 with no overall bias being present (i.e., no large difference between seeded and nonseeded periods).

As an initial examination of the data, the monthly rainfall is averaged over the seeded and nonseeded months (Table 2). Looking at the 50% longevity filter on the data, both the HECA and CSIRO targets find roughly a 20% increase in rainfall during periods of cloud seeding relative to nonseeded periods. The assumption that this increase is due to seeding may be misleading, as the differences may simply reflect decadal variations in rainfall in the target area. The double ratio (Gabriel 1999) has been a common analysis tool employed in cloud seeding research to account for temporal variations in rainfall. To perform a double ratio a control region must be defined as having an area-average monthly rainfall during nonseeded periods that is highly correlated with the target (either the HECA or CSIRO) and that is not affected by the act of seeding. As the primary unit of time is one calendar month, it is not possible to say that the control areas used are not affected by the seeding, as over the course of a month it is entirely possible to have winds that span the full 360° (Fig. 3). Indeed, the act of seeding 60 min upwind of the target, as was done in Tasmania 2, often led to the seeding tracks lying off the west coast of Tasmania.

Correlation maps of monthly rainfall totals for individual rain gauges with respect to both the HECA and CSIRO targets during nonseeded months (Fig. 6) are used to define the controls. Excluding the increased correlation observed for the southern portion of the HECA, little difference is observed in the overall correlation pattern regarding either the HECA or CSIRO target. In defining control regions within Tasmania, only stations outside of the HECA target were considered. In an effort to minimize any potential extra area effects of the seeding, an ~25-km buffer was added around the HECA target. Controls were only defined outside of this buffer region. The correlation map may roughly be broken up into four broad geographic regions: west, northwest, northeast, and southeast (hereinafter W, NW, NE, SE control regions) of the target. In addition to these four geographically based controls, a fifth control has been defined using 10 ''high-quality'' (HQ) rain gauges developed by Lavery et al. (1997; Fig. 7). None of these 10 sites resides within the HECA or CSIRO target although one of the sites does reside within the 25-km buffer zone to the southeast of the target. The nine remaining HQ sites are found within the four regional controls as follows: W (one), NW (three), NE (three), and SE (two).

The surface rainfall observations for the five different controls were prepared in a manner identical to that for the two targets. Similar to Table 2, the monthly rainfall for the five controls may be averaged over seeded and nonseeded months (Table 3). Only the 50% longevity threshold data are presented. Whereas both the HECA and CSIRO targets realized a roughly 20% gain in rainfall during seeded periods, the controls for the W, NW, NE, and SE display gains of 14.1%, 10.5%, 5.1%, and 6.5%, respectively. These results show that rainfall

TABLE 2. Average rainfall (mm) for all months (May–October 1960–2005), only seeded months, nonseeded months, and lightly seeded months (1–2 seeds) for both the CSIRO and HECA targets for the 25%, 50%, and 75% longevity filters.

	No. of months	CSIRO 25%	CSIRO 50%	CSIRO 75%	HECA 25%	HECA 50%	HECA 75%
All months	276	118.0	113.9	127.0	162.6	160.2	171.3
Seeded months	130	131.7	124.7	139.4	179.4	175.9	185.6
Nonseeded months	146	105.9	104.3	115.9	147.6	146.3	158.6
Lightly seeded months	60	117.7	112.5	125.0	159.6	157.0	164.6

is greater during seeded periods over the whole of Tasmania and suggest that underlying temporal variations in rainfall should not be neglected.

Using only the 146 nonseeded months, the correlation between each control with either target is presented in Table 3. Based strictly on the geometry of the control regions, it is not surprising that the W control demonstrates a higher correlation than the other regional controls. The NW, NE, and SE controls are larger in size and contain more sites at a farther distance from the target. The SE control is not particularly well correlated with either target in comparison to the W and NW. When swapping from the larger HECA target to the smaller CSIRO target, the correlation with the W control drops from 0.87 to 0.79, while the correlation increases for the NE and SE controls. This is consistent with the geometry of the controls. The CSIRO target is not as mountainous as the HECA and therefore shows a higher correlation with the relatively low-lying NE and SE controls. The HECA, being relatively more mountainous, exhibits a higher correlation with other like regions, those being the W and NW.

One might expect that the HQ control would roughly act as a weighted average of the four regional controls. This is not the case, as the HQ control displays only a weak enhancement during the seeded months (5.8%) but maintains a high correlation with both the larger HECA target (0.85) and the smaller CSIRO target (0.87).

5. Double ratio analysis

Given that the four regional controls find rainfall enhancements from 5% to 14% (Table 3) during seeded months, the increase of 20% during seeded months for the two targets (Table 2) cannot be taken at face value. The conventional means of removing temporal variations in rainfall is to use the double ratio (Gabriel 1999). If the target rainfall for month i is T_i and the control rainfall is C_i and the summation over all seeded (unseeded) months is Σ_s (Σ_u), then the double ratio d is defined as

$$d = \frac{\sum_s T_i \Big/ \sum_s C_i}{\sum_u T_i \Big/ \sum_u C_i}, \tag{1}$$

or alternatively,

$$1.0 = \frac{\sum_s (T_i/d) \Big/ \sum_s C_i}{\sum_u T_i \Big/ \sum_u C_i}. \tag{2}$$

The 10 possible double ratios (Fig. 8) range from as little as 1.047 (CSIRO versus W) to as large as 1.145 (HECA versus NE). These values, while considerably less than the 20% found in Table 2, are still noteworthy. The double ratio suggests that over the 46-yr period, cloud seeding increased rainfall in the target between 5% and 14% relative to nearby controls. Following the method of Shaw et al. (1984), 95% confidence intervals may be defined for these double ratios. The estimated increase associated with a double ratio d is

$$\left(1 - \frac{1}{d}\right) \sum_s T_i = \sum_s T_i - \sum_u T_i \frac{\sum_s C_i}{\sum_u C_i}. \tag{3}$$

If (d_1, d_2) is a $100(1 - \alpha)\%$ confidence interval for d, then $\{[1 - (1/d_1)], [1 - (1/d_2)]\}$ is the $100(1 - \alpha)\%$ confidence interval for the increase. Suppose that in Eq. (2), instead of dividing T_i by d, T_i is divided by some δ; any double ratio calculated is therefore $d(\delta)$, with $d(\delta) < 1$ if $\delta > d$ and $d(\delta) > 1$ if $\delta < d$. It is therefore possible to associate a permutation estimate of the probability of a $d(\delta)$ different from 1.0, here defined as $p(\delta)$. The required confidence limits d_1 and d_2 are therefore the two solutions of $p(\delta) = \alpha/2$ and are found numerically. In essence, two questions are asked. First, at what value for δ does the probability of obtaining a double ratio $d(\delta)$ greater than 1 exceed 97.5%? Second, at what value of δ does the probability of obtaining a double ratio $d(\delta)$ greater than 1 drop below 2.5%? The value of δ is found by trial and error and the probabilities are assessed using a standard bootstrap technique (Efron and Tibshirani 1993).

It is observed that all of the confidence intervals lie above a value of one when using the HECA target (Fig. 8a), and all overlap between about 6% and 11%. It is interesting to note that while the HECA–HQ double ratio is comparable to the HECA–SE and HECA–NE values, the confidence interval is smaller. This reflects the higher correlation that the HECA target has with the HQ control over the NE and SE controls. Appropriately, the confidence interval broadens as the correlation drops. The confidence intervals when employing the CSIRO target (Fig. 8b) are comparable with the exception that the 95% confidence interval for the CSIRO-W does extend below 1.00.

While these confidence intervals suggest that the positive response to cloud seeding is physically consistent across all regions, they do not rigorously define the statistical significance of the result. This is because the confidence intervals address the uncertainty in the calculated double ratio. They cannot address the question, how likely is it that this double ratio could occur by

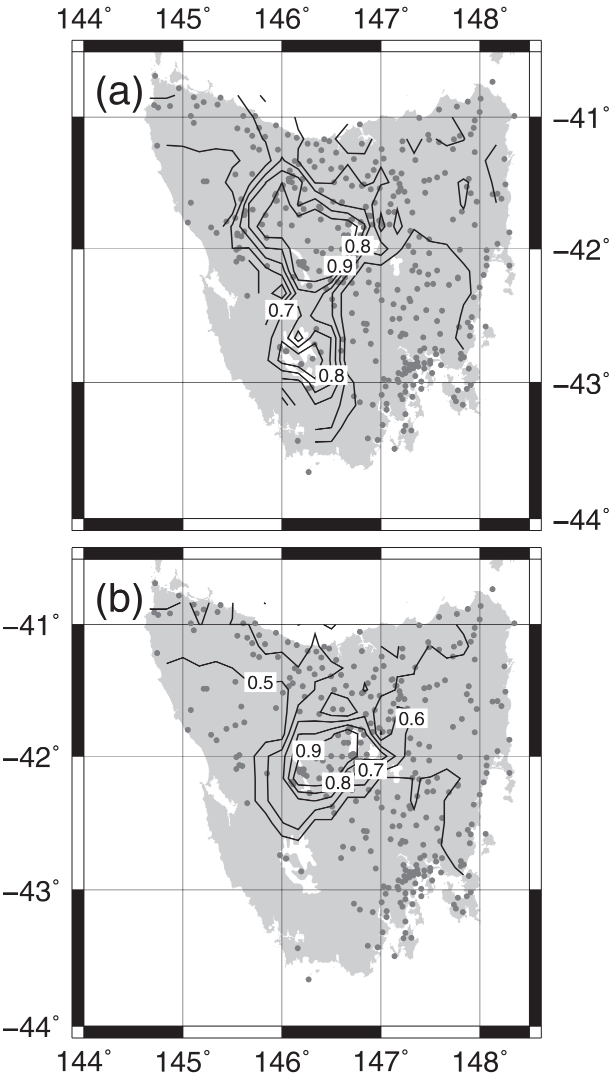

FIG. 6. Correlation maps for (a) HECA and (b) CSIRO target areas for the 50% threshold filter. Correlation is for unseeded periods only. Regions containing highly correlated sites closely resemble target perimeter.

purely random processes? A bootstrap analysis (Efron and Tibshirani 1993) is undertaken for this purpose. For this analysis, 10 000 bootstrap samples were calculated for the double ratio. Each bootstrap sample is constructed by randomly drawing (with replacement) a month from the full pool of 276 months 130 times. The target and control pair from these 130 selected months is defined as "seeded." Similarly, 146 "nonseeded" months are defined by 146 independent draws from the full data pool of 276 (i.e., replacement is enforced). This bootstrap sample double ratio is then compared with the true double ratio. The number of times that a bootstrap sample exceeds the true value is simply counted to define the significance level. If 500 of the 10 000 bootstrap samples exceed the true double ratio, then it is stated that the gain is at the 95% significance level. Of the 10 combinations possible (Fig. 8), only one fails to reach the 95% level (the CSIRO-W is at 92.9%). In the case of the HECA target, all five comparisons surpass the 98% level. Both the HECA–HQ and CSIRO–HQ double ratios surpass the 99.99% level of significance. As there is more than one comparison being employed, the chances of obtaining a statistically significant result of a double ratio being different from 1.0 are greatly increased. This is because multiple statistical inferences are being considered simultaneously. Applying the Bonferroni adjustment (Weisstein 2008) for multiple comparisons to a 95% significance level for 10 tests produces an alpha of 0.005. In this test, 6 out of the 10 comparisons pass at the required level of significance.

These high levels of significance must ultimately be tempered since, strictly speaking, the monthly rainfall does exhibit both a month-to-month and year-to-year autocorrelation. The bootstrap analysis assumes that all 276 single ratios are independent of one another. Ideally, the monthly cloud seeding would have taken place completely randomly across the 6 months and 46 yr. Operationally this is simply not practical. The month-to-month seasonally adjusted autocorrelation at lag one over the full 276 months for the HECA (CSIRO) target is 0.129 (0.105), after transforming to a standard normal distribution. This is comparable to the month-to-month autocorrelation found over the full 12-month, 46-yr dataset. The year-to-year autocorrelation at lag 1 was calculated to be -0.043 and -0.140 for the HECA and CSIRO targets, respectively (once again, after transforming to a standard normal distribution). It is noted that autocorrelations less than 0.2 are generally not considered particularly strong.

It is possible to eliminate any month-to-month correlation by analyzing the months individually; instead of one dataset of 276 months, six separate datasets of 46 months may be tested. This approach has two drawbacks. First, the concern about multiple comparisons is magnified. Instead of 10 tests there are now 60, the two target areas compared against the six months and five control regions. Second, the reduced sample size makes it much more difficult to reach statistical significance. The results of this analysis remain consistent with those presented thus far. However, because of reduced sample size and the issue regarding multiple comparisons, the results show much greater variability and are considerably less significant.

6. Discussion

An analysis of the surface rainfall data over Tasmania for the period of 1960–2005 readily finds that, on

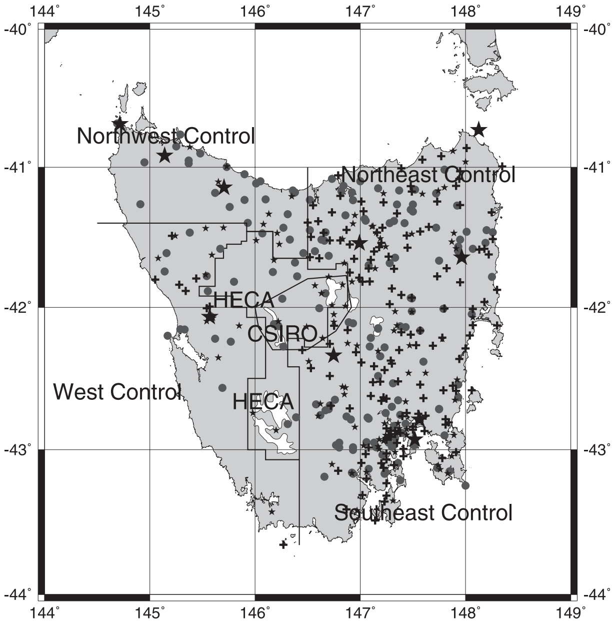

FIG. 7. Map of Tasmania highlighting the locations of sites that pass the 25%, 50%, and 75% longevity filters (small gray circles, small black stars, and small black crosses, respectively) for the period 1960–2005. The HECA and CSIRO target areas together with the control areas are defined (note that any sites inside a control area within 25 km of the target perimeter are not included in the control region). Sites marked using a large black star are Bureau of Meteorology HQ sites.

average, more rainfall did occur during months in which seeding took place within target and control regions. The intermittent nature of the seeding (both experimental and operational) over this extended period provides a unique dataset that may ultimately help demonstrate that cloud seeding can be viable over an economically meaningful area.

A standard double ratio calculation finds that the rainfall over the target was somewhere between 5% and 13% greater than over nearby ''control'' regions with a satisfactory level of statistical significance being reached (using a bootstrap analysis) in many of the tests. Furthermore, it is thought that the consistency in the findings lends much to the credibility. Both Tasmania 1 and 2 present analyses indicating an increase in precipitation associated with seeding; the present analysis is consistent with these.

TABLE 3. Correlations between control areas and target areas. Also shown is the average rainfall (May–October 1960–2005), average seeded rainfall, and average unseeded rainfall (mm) for each control.

	W	NW	NE	SE	HQ
Correlation with HECA	0.87	0.76	0.66	0.51	0.85
Correlation with CSIRO	0.79	0.74	0.73	0.64	0.87
Avg rainfall	223.11	127.19	84.01	63.29	95.51
Seeded rainfall	238.76	133.91	86.21	65.40	98.38
Nonseeded rainfall	209.17	121.21	82.05	61.41	92.96

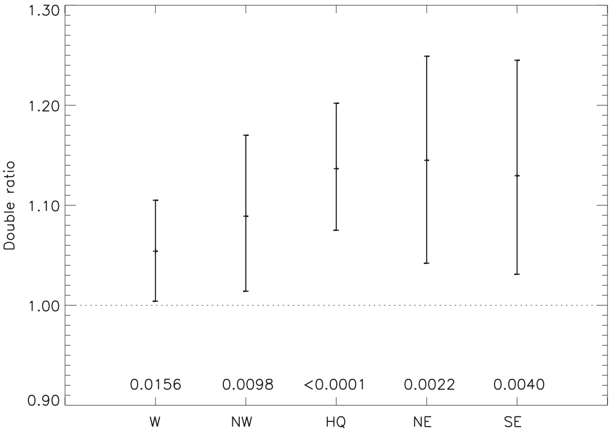

(a) Results for the pooled 276 month data set for the HECA target.

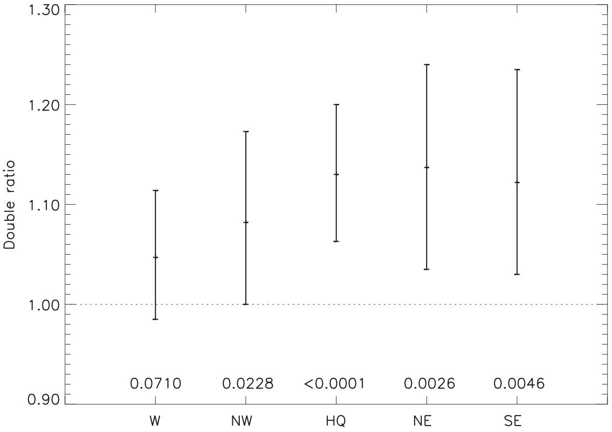

(b) Results for the pooled 276 month data set for the CSIRO target.

FIG. 8. Double ratios and confidence intervals for the HECA and CSIRO targets vs the W, NW, HQ, NE, and SE controls. The bootstrap probabilities for obtaining a double ratio higher than the actual are shown above the horizontal axis.

As the cloud seeding projects over Tasmania were never designed as a single long-term field experiment, numerous caveats to this approach exist. Evolving target boundaries, changing seeding strategies and technologies, and evolving surface sites all limit the findings. More importantly, the seeding was not undertaken randomly on a monthly basis. Month-to-month and year-to-year correlations exist within the rainfall observations, which although small, violate assumptions of the bootstrap analysis.

Another major caveat to this approach is in the definition of control areas from which to define a double ratio. Control areas are supposed to be highly correlated with the target area during nonseeded periods, but free of any immediate effect of cloud seeding. When examining the data on a monthly time scale, there is no location within Tasmania that is not either downwind of the target or upwind where seeding commonly occurs. A 25-km buffer zone was enforced around the larger target in an effort to minimize any potential extra-area effect of seeding. It is interesting to note that the closest control with the highest chance of extra-area effects consistently shows the lowest double ratio with the least significance.

Furthermore, the effects of both wind shear and frontal deformation on the targeting and dispersion of the seeding agent remain very legitimate questions. Any extra-area effects would ultimately serve to diminish the double ratio and underestimate the effectiveness of the cloud seeding. It is also noted that operational constraints during seeding periods such as aircraft/instrument maintenance, personnel limitations, inaccurate forecasting, and restrictions near lightning also lead to an underestimation of the potential efficiency.

Bearing in mind these caveats, this analysis provides some support for the hypothesis that cloud seeding may be physically plausible over Tasmania. This analysis cannot, however, provide ‘‘convincing scientific proof,’’ even when coupled with the research of the Tasmania 1 and 2 experiments. Many of the key uncertainties identified in the U.S. NRC report still remain. Most notably, the basic microphysical state of these precipitating frontal systems over the Southern Ocean is essentially unknown. These are not wintertime orographic clouds, as denoted by the NRC; these cloud systems reside within the high wind shear ‘‘roaring forties.’’ Moreover, the atmosphere over the Southern Ocean is pristine with little terrestrial or anthropogenic influence evident on the time scale from days to weeks. Recent observations suggest that even boundary layer clouds over the Southern Ocean are notably different from pristine maritime boundary layer clouds observed in the Northern Hemisphere (Bennartz 2007). The conclusions reached from glaciogenic seeding efforts over the western United States (Deshler and Reynolds 1990; Deshler et al. 1990) are unlikely to be applicable over Tasmania.

The cloud systems over the Southern Ocean cover over 10% of the earth's surface yet remain largely unexplored. The unique dynamics and microphysics of these systems present a particular challenge in understanding the targeting, dispersion, and physical response of the clouds to glaciogenic cloud seeding. Obviously such challenges may only be addressed through further fieldwork. If any such fieldwork were to establish the immediate effect of cloud seeding, then it might be argued that, together with the fieldwork of Tasmania 1 (Smith et al. 1979), 2 (Ryan and King 1997), and the present analysis, a complete argument for the efficacy of cloud seeding has been presented.

Acknowledgments. This research has been supported by the Australian Research Council through Linkage Grant LP0562358. The authors are particularly indebted to Ian Searle, Christina Nebel, and Mark De Hoog. Adian Sudbury and Kais Hamza are thanked for their help in the statistical analysis. The Tasmanian regional office of the Bureau of Meteorology and Blair Trewin are thanked for their assistance regarding the field data and controls. Roelof Bruintjes, William Cooper, and Tara Jensen are thanked for timely advice. The manuscript has been revised heavily in the review process, for which we are also grateful. Last, the authors thank the late Brian Ryan for inspiring much of this research.

REFERENCES

Baumgardner, D., H. Jonsson, W. Dawson, D. O'Conner, and R. Newton, 2001: The cloud, aerosol and precipitation spectrometer: A new instrument for cloud investigations. *Atmos. Res.,* **59–60,** 251–264.

Bennartz, R., 2007: Global assessment of marine boundary layer cloud droplet number concentration from satellite. *J. Geophys. Res.,* **112,** D02201, doi:10.1029/2006JD007547.

Boe, B., G. W. Bomar, W. R. Cotton, B. L. Marler, H. D. Orville, and J. A. Warburton, 2004: The Weather Modification Association's response to the National Research Council's report titled: ‘‘Critical issues in weather modification research.’’ *J. Wea. Modif.,* **36,** 53–82.

Deshler, T., and D. W. Reynolds, 1990: The persistence of seeding effects in a winter orographic cloud seeded with silver iodide burned in acetone. *J. Appl. Meteor.,* **29,** 477–488.

——, ——, and A. W. Huggins, 1990: Physical response of winter orographic clouds over the Sierra Nevada to airborne seeding using dry ice or silver iodide. *J. Appl. Meteor.,* **29,** 288–330.

Efron, B., and R. J. Tibshirani, 1993: *An Introduction to the Bootstrap. Monogr. Stat. Appl. Probab*., No. 57, Chapman and Hall/CRC, 436 pp.

Gabriel, K. R., 1999: Ratio statistics for randomized experiments in precipitation stimulation. *J. Appl. Meteor.,* **38,** 290–301.

——, and D. Rosenfeld, 1990: The second Israeli rainfall stimulation experiment: Analysis of precipitation on both targets. *J. Appl. Meteor.,* **29,** 1055–1067.

Garstang, M., R. Bruintjes, R. Serafin, H. Orville, B. Boe, W. Cotton, and J. Warburton, 2005: Finding common ground. *Bull. Amer. Meteor. Soc.,* **86,** 647–655.

Gras, J. L., 1995: CN, CCN and particle size in Southern Ocean air at Cape Grim. *Atmos. Res.,* **35,** 233–251.

Jarque, C. M., and A. K. Bera, 1980: Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Econ. Lett.,* **6,** 255–259.

Kraus, E. B., and P. Squires, 1947: Experiments on the stimulation of clouds to produce rain. *Science,* **159,** 489–491.

Lavery, B., G. Joung, and N. Nicholls, 1997: An extended high-quality historical rainfall dataset for Australia. *Aust. Meteor. Mag.,* **46,** 27–38.

Long, A. B., and A. W. Huggins, 1992: Australian Winter Storms Experiment (AWSE) I: Supercooled liquid water and precipitation-enhancement opportunities. *J. Appl. Meteor.,* **31,** 1041–1055.

Mace, G. G., R. Marchand, Q. Zhang, and G. Stephens, 2007: Global hydrometeor occurrence as observed by CloudSat: Initial observations from summer 2006. *Geophys. Res. Lett.,* **34,** L09808, doi:10.1029/2006GL029017.

McGregor, J. L., and M. R. Dix, 2007: An updated description of the Conformal-Cubic Atmospheric Model. *High Resolution Numerical Modelling of the Atmosphere and Ocean*, K. Hamilton and W. Ohfuchi, Eds., Springer, 51–75.

National Climate Centre, 2006: An exceptionally dry decade in parts of southern and eastern Australia: October 1996–September 2006. Australian Bureau of Meteorology Special Climate Statement 9, 9 pp. [Available online at http://www.bom.gov.au/climate/current/statements/scs9a.pdf.]

Nirel, R., and D. Rosenfeld, 1995: Estimation of the effect of operational seeding on rain amounts in Israel. *J. Appl. Meteor.,* **34,** 2220–2229.

Rauber, R. M., and L. O. Grant, 1987: Supercooled liquid water structure of a shallow orographic cloud system in southern Utah. *J. Climate Appl. Meteor.,* **26,** 208–215.

Reynolds, D. W., and A. S. Dennis, 1986: A review of the Sierra Cooperative Pilot Project. *Bull. Amer. Meteor. Soc.,* **67,** 513–523.

Rosenfeld, D., and W. L. Woodley, 1989: Effects of cloud seeding in west Texas. *J. Appl. Meteor.,* **28,** 1050–1080.

Ryan, B. F., and W. D. King, 1997: A critical review of the Australian experience in cloud seeding. *Bull. Amer. Meteor. Soc.,* **78,** 239–254.

Schaefer, V. J., 1946: The production of ice crystals in a cloud of supercooled water droplets. *Science,* **104,** 457–459.

Shaw, D. E., W. D. King, and E. Turton, 1984: Analysis of Hydro-Electric Commission cloud seeding in Tasmania 1979–83. CSIRO Division of Cloud Physics CP Tech. Rep. 394, 33 pp.

Silverman, B. A., and W. Sukarnjanasat, 2000: Results of the Thailand warm-cloud hygroscopic particle seeding experiment. *J. Appl. Meteor.,* **39,** 1160–1175.

Smith, E. J., E. E. Adderley, and D. Walsh, 1963: A cloud-seeding experiment in the Snowy Mountains. *J. Appl. Meteor.,* **2,** 324–332.

——, L. G. Veitch, D. E. Shaw, and A. J. Miller, 1979: A cloud-seeding experiment in Tasmania. *J. Appl. Meteor.,* **18,** 804–815.

Watkins, A., and B. Trewin, 2007: Australian climate summary: 2006. *Bull. Aust. Meteor. Oceanogr. Soc.,* **20,** 10–17.

Weisstein, W. E., cited 2008: Bonferroni correction. Wolfram MathWorld. [Available online at http://mathworld.wolfram.com/BonferroniCorrection.html.]