A Confirmatory Snowfall Enhancement Project in the Snowy Mountains of Australia. Part II: Primary and Associated Analyses

MICHAEL J. MANTON

Monash University, Melbourne, Victoria, Australia

LOREDANA WARREN

Snowy Hydro Limited, Sydney, New South Wales, Australia

(Manuscript received 20 September 2010, in final form 9 February 2011)

ABSTRACT

The Snowy Precipitation Enhancement Research Project (SPERP) was undertaken in winters from May 2005 to June 2009 in the Snowy Mountains region of southeastern Australia. Part I of this paper describes the design and implementation of the project, as well as the characteristics of the key datasets collected during the field phase. The primary analysis in this paper (Part II) shows an unequivocal impact on the targeting of seeding material, with the maximum level of silver in snow samples collected from the primary target area found to be significantly greater in seeded than unseeded experimental units (EUs). A positive but not statistically significant impact on precipitation was found. Further analysis shows that a substantial source of uncertainty in the estimation of the impacts of seeding on precipitation is associated with EUs where the seeding generators operated for relatively few hours. When the analysis is repeated using only EUs with more than 45 generator hours, the increase in precipitation in the primary target area is 14% at the 8% significance level. When applying that analysis to the overall target area, the precipitation increase is 14% at the 3% significance level. A secondary analysis of the ratio of silver to indium in snow supports the hypothesis that seeding material affected the cloud microphysics. Other secondary analyses reveal that seeding had an impact on virtually all of the physical variables examined in a manner consistent with the seeding hypothesis.

1. Introduction

A confirmatory cloud seeding experiment has been carried out in the Snowy Mountains of southeastern Australia, using silver iodide dispersed from ground-based generators to induce the transformation of supercooled liquid water into ice during westerly flows associated with the passage of winter cold fronts. The seeding hypothesis for the Snowy Mountains Precipitation Enhancement Project (SPERP) is that the silver iodide acts as ice nuclei, which transform the supercooled liquid water into ice particles that fall as snow in the fixed target area along the mountain ranges. Precipitation in a control area to the west and north of the target is used to estimate the natural precipitation in the target. The primary target area is at elevations greater than about 1400 m, and it is expected that induced precipitation will generally fall as snow in that area (see Fig. 2 of Manton et al. 2011). A secondary target extends to the east and west of the primary target and it includes ranges at lower elevations. Together the primary and secondary targets form the overall target with an area of 832 km^2.

Manton et al. (2011, hereinafter Part I) describes the design and implementation of SPERP, which ran from May 2005 to June 2009 during the winter months. During the experiment, 107 five-hour experimental units (EUs) were declared when specified conditions for seeding were satisfied. Seeding was randomized with a seeding ratio of 2:1. A passive tracer [indium (III) oxide] was dispersed during every EU. In this paper, the results of the primary analysis of SPERP are presented and discussed.

2. Primary analysis

The primary analysis comprises two key components (Part I). The first component is the detection of

Corresponding author address: Michael Manton, School of Mathematical Sciences, Monash University, Melbourne, VIC 3800, Australia.
E-mail: michael.manton@monash.edu

DOI: 10.1175/2011JAMC2660.1

a seeding impact on precipitation based on a regression analysis in which precipitation in the control area in unseeded EUs is used to estimate the "natural" precipitation in the primary target area. The second component is the confirmation that seeding material has reached the primary target area, based on measurements of the maximum value of silver (Ag) in seeded and unseeded EUs.

The experiment was designed to have a seeding ratio of 2:1; that is, the number of seeded EUs should be twice the number of unseeded EUs. Of the 107 EUs, 71 were seeded and 36 were unseeded. As would be expected, the number of EUs each year was affected by the weather, and there were reduced numbers in 2006 and 2009, which were weak El Niño years [when precipitation in eastern Australia can be below average (Nicholls 1989)]; moreover, operations for this experiment in 2009 occurred only in June.

Inspection of the month-by-month distribution of EUs shows that the number of EUs in May and September is less than the numbers in June–August. The opportunities for EUs in autumn and spring are especially limited by the legislated condition that seeding can be conducted only when snow is falling. Recognizing the large interannual variability of seeding opportunities in the region (Manton et al. 2009), the actual distribution of EUs is within the historical range. Indeed, the design of SPERP was based on the assumption that about 100 EUs could be observed over a 5-yr period.

a. *Impacts of seeding on precipitation in primary target*

Using the core dataset for the primary analysis and applying the seeding sequence to the precipitation data, we find that the seeded events tend to have higher precipitation than unseeded events for both the target and control areas. The median precipitation in the primary target area is 2.65 mm for seeded EUs and 2.08 mm for unseeded. In the control area, the median values are 1.90 mm for seeded and 1.00 mm for unseeded EUs.

The maximum control precipitation in unseeded EUs is 8.73 mm compared with 11.82 mm for seeded EUs. Ideally, the distribution of precipitation in the control area would be the same for seeded and unseeded EUs, so that the regression analysis spans the same range as that to which it is applied for prediction. However, a Kolmogorov–Smirnov (KS) test (Marsaglia et al. 2003) shows that the distribution of precipitation in the control area is significantly different (at the 8% level) in seeded and unseeded EUs. A possible reason for the precipitation range in the control area being smaller in unseeded than in seeded EUs is the use of a seeding ratio greater than 1:1; that is, there are fewer opportunities for very large values to occur in the unseeded cases.

Although there are many EUs with relatively low precipitation values, the core precipitation dataset should provide an adequate basis for the analysis of the impacts of seeding, bearing in mind a number of issues that could lead to uncertainties in the results of the primary analysis. A linear regression is applied to these data in the unseeded EUs (Part I), so that the natural precipitation in the target area can be predicted from the precipitation in the control area. If TU and CU are the precipitation in the target and control areas for the unseeded EUs, then we can compute the regression coefficients, *a* and *b*, where

$$\mathrm{TU} \sim a + b \times \mathrm{CU}. \tag{1}$$

The computed regression coefficients are then used in the seeded EUs to estimate the natural precipitation in the target area in the absence of any seeding impacts. The difference between the observed precipitation in the target area and the estimated natural precipitation is the seeding impact for each EU. If TS and CS are the precipitation in the target and control areas for seeded EUs, then the seeding impact RS is given by

$$\mathrm{RS} = \mathrm{TS} - (a + b \times \mathrm{CS}). \tag{2}$$

The variable RS is the regression residual, as it represents the difference between the actual and estimated values of precipitation in the target area. The specified indicator of the overall seeding impact is the total fractional increase (TFI) in precipitation above the natural level in the target area, where

$$\mathrm{TFI} = \mathrm{sum(RS)}/[\mathrm{sum(TS)} - \mathrm{sum(RS)}].$$

It is expected that TFI should be greater than zero; that is, the sum of the residuals is positive. To confirm the estimated value, it is necessary to show that the value of TFI is greater than what could occur by chance at the 10% level of significance. The relatively low significance level of 10% is specified in Part I on the basis of simulations of the probability of detection of a seeding impact using historical precipitation data. (The higher the percentage value, the greater the probability that the event could occur by chance; so a low significance level is associated with high percentage values.)

A bootstrap method (Part I) is used to determine the significance level of the observed value of TFI. From the 107 EUs, 71 are randomly selected (with replacement) as "seeded" and 36 as "unseeded." The regression analysis is used to estimate *a* and *b* in Eq. (1), and then the residual for seeded EUs is computed from Eq. (2). The residuals are summed and normalized to yield the TFI for that simulation. The simulation is repeated 100 000 times to determine the distribution of TFI values that are possible from the SPERP dataset. The one-sided significance of the TFI is estimated by the fraction of the distribution having values greater than the observed value of TFI.

In addition to estimating the statistical significance of the computed value of TFI, a second bootstrap method is used to determine the confidence interval for the actual value of TFI; in particular, 71 seeded and 36 unseeded samples from the 36 unseeded EUs are chosen at random (with replacement) 100 000 times with the levels of the selected seeded cases increased by the observed seeding impact (TFI). From the overall distribution of simulated values of TFI, we find the 2.5% and 97.5% quantiles to estimate the likely range of TFI at the 95% level.

From the unseeded EUs, the regression coefficients, *a* and *b*, are found to be 0.83 mm and 1.10 respectively. Figure 1 shows a scatterplot of the target and control precipitation amounts, together with the regression line derived from the unseeded EUs. This diagram highlights the relative absence of unseeded EUs with precipitation in the control area covering the upper range of values in seeded EUs. Thus, while it appears that many of the seeded EUs have precipitation above the regression line, there must be some uncertainty about the validity of the regression at large values of the control precipitation. Moreover, inspection of Fig. 1 suggests that there may also be uncertainties at small values of the control precipitation where the regression line tends to overestimate the target value.

We find that the TFI is calculated to be 0.07; that is, the regression analysis implies that seeded EUs yield 7% more precipitation in the primary target than unseeded EUs. However, the bootstrap analysis suggests that there is a 24% probability that this value of the fractional increase could occur by chance; that is, the value does not reach the required significance level of 10%. The second bootstrap analysis implies that the confidence interval for the fractional increase is from −0.16 to 0.32, so a possible range of the seeding impact is from a 16% reduction to a 32% increase in precipitation.

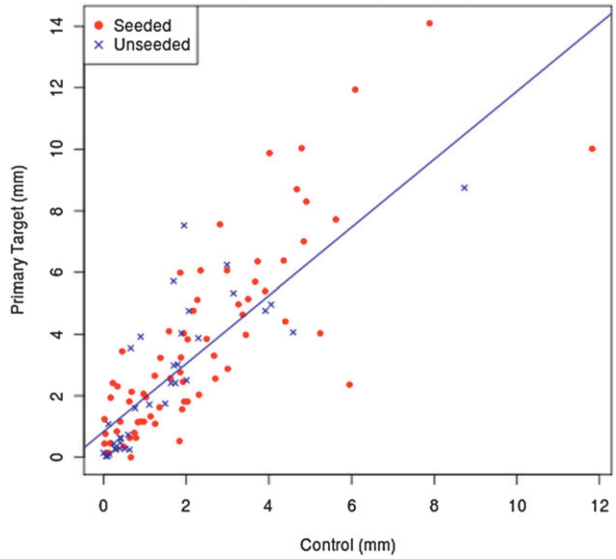

FIG. 1. Scatterplot of precipitation in primary target and control areas; the blue line denotes the regression line based on unseeded EUs.

b. *Impacts of seeding on precipitation in overall target*

As shown in section 4a of Part I, the correlation in precipitation between the control and overall target (primary plus secondary) areas is higher (0.90) than for the primary target area alone (0.82), which leads to greater expectations that a seeding impact can be detected. It is therefore appropriate to consider the impacts of seeding on the overall target area.

Repeating the primary analysis using the precipitation in the overall target area, the TFI is found to be 9% at the 13% significance level, and the confidence interval is −0.13 to 0.31. Not only is the seeding impact higher in the overall target (9% compared with 7%), but the statistical significance is substantially improved at 13%. (The statistical result would be stronger if the confidence interval was entirely positive.)

The effects of the higher correlation between the overall target and control compared with primary target and control are greater in SPERP than in the historical analysis (Part I), where the correlation increases from 0.71 to 0.79. In fact, the primary target was selected on the basis that it is the region where any enhanced precipitation must fall as snow; that is, it is above about 1400-m elevation. The secondary target area is lower in elevation, but nonetheless it is expected to be targeted equally well as the primary target (Part I). Recognizing the results of the formal primary analysis, the focus of subsequent analysis will be on the overall (rather than primary) target area.

c. *Targeting of seeding material*

The evaluation plan for SPERP specifies that the snow chemistry variable for the primary analysis (Part I) is the peak value of Ag over sites in the primary target area where snow-sample observations can be unambiguously aligned with an EU and where the concentration of indium (In) is larger than 1 ppt. The required result is achieved if the peak value of Ag (MaxAg) is on average larger during seeded EUs than during unseeded EUs. A Wilcoxon rank-sum test (Bauer 1972) is used to demonstrate that the difference is significant at the 5% level.

The core dataset is obtained by combining the chemistry data from the profiles and from the Blue Cow site in the primary target area. There are 441 quality-controlled snow slices that have In > 1 ppt and that can be assigned unambiguously to one EU; 185 of the slices are from Blue Cow and 256 from the profiles. Figure 2 shows the value of MaxAg in the primary target area for each EU. Noting that there are no snow chemistry data for 25 EUs, we find that there are 56 seeded EUs and 26 unseeded EUs that can be assigned a value of MaxAg. Thus, the 2:1 seeding ratio is essentially maintained in the chemistry data. The MaxAg in seeded EUs ranges from 4.16 to 147 ppt with a median value of 31.7 ppt. For unseeded EUs, the range is 0.49 to 40.4 ppt with a median of 9.59 ppt. Inspection of Fig. 2 shows that the values of MaxAg in unseeded EUs are consistently lower than those in the seeded EUs. This subjective assessment is formally confirmed by the Wilcoxon test, which yields a p value of 1×10^{-6}; that is, the difference in the distributions is very significant. It is apparent from Fig. 2 that, owing to uncertainties in the timing of the snow chemistry samples, dispersion along a snow profile, or wind-blown snow between EUs (Part I), there may be some correlation between unseeded EUs and adjacent seeded EUs during field campaigns.

The requirement of the primary analysis for the targeting of seeding material is therefore readily met: the MaxAg in seeded EUs is larger than that in unseeded EUs with a statistical significance of at least 5%. It can be argued that the consistent presence of samples with In > 1 ppt itself indicates that targeting has been successful (Part I). Given the range of values of Ag in unseeded EUs (the median concentration of Ag in unseeded samples is 5.4 ppt), however, it is reassuring to confirm that Ag is substantially higher in seeded compared to unseeded samples.

d. *Ratio of silver to indium in primary target*

Having readily satisfied the primary analysis test for targeting, it is appropriate to consider the robustness of the earlier results of Chai et al. (1993) on the ratio of Ag:In when there is a rigorous control on the timing of the snow samples. The dataset used in section 2c comprises 340 samples from seeded EUs and 101 from unseeded EUs, with In greater than 1 ppt in all cases. The ratio of Ag:In for unseeded samples varies from 0 to 12.3 with a median of 0.88, while Ag:In varies from 0.31 to 26.5 for seeded EUs with a median of 2.99. Figure 3

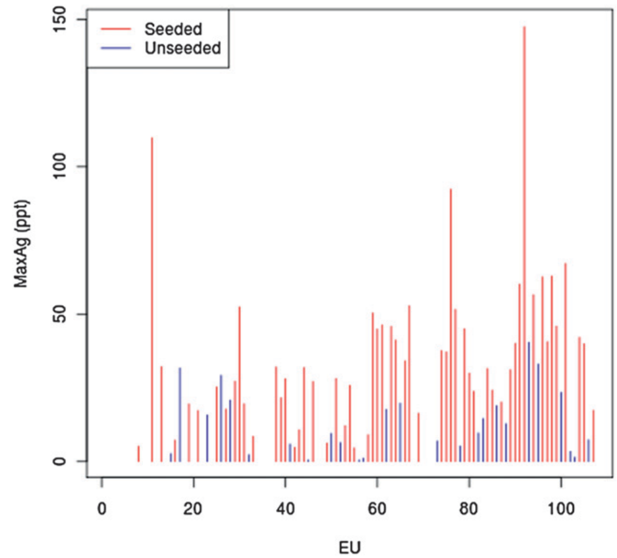

FIG. 2. MaxAg in the primary target area for each EU; red denotes seeded EUs and blue unseeded EUs.

shows a scatterplot of In and Ag for seeded and unseeded cases, with a line at the median slope of Ag:In for seeded and unseeded cases. A KS test confirms that the distributions are very significantly different. A value of one for Ag:In for seeded EUs would imply that Ag (as for In) was being removed only by scavenging.

The results of the analyses in sections 2c and 2d support the hypothesis that seeding material is reaching the target area and is influencing the precipitation process during seeded EUs (rather than just being scavenged). The results of Chai et al. (1993) are therefore statistically supported when there is a rigorous separation of seeded and unseeded events.

Although the results of the snow chemistry analysis are encouraging, it must be noted that large concentrations of Ag are sometimes found in unseeded events; in particular, the median value is 5.4 ppt with a maximum of 40 ppt. Huggins et al. (2008) suggest that the background levels of Ag and In are less than 3 ppt. The high values of Ag may simply represent contamination due to wind-blown snow between EUs, but further work is needed to clarify the nature and causes of the actual distribution of Ag.

3. Sensitivity of seeding impacts to amount of seeding material

The maximum coverage over the target is 65 generator hours per EU (Part I). However, as found earlier, there was a wide variation in the generator hours, with about 20% of EUs having less than 45 h. As it is

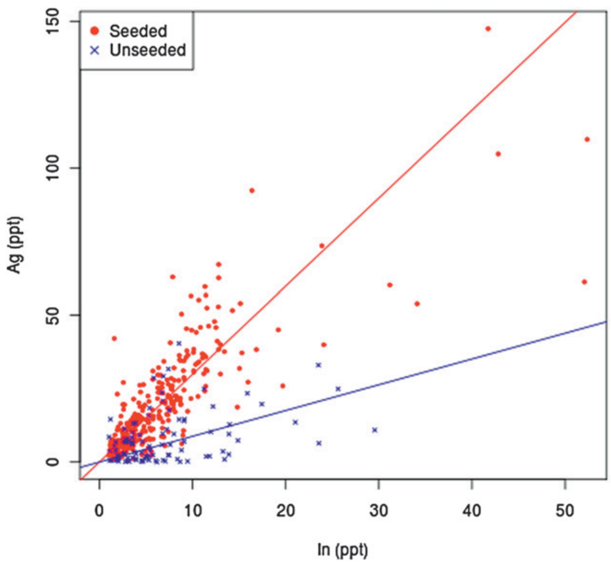

FIG. 3. Scatterplot of concentrations of In and Ag in snow samples in seeded (red) and unseeded (blue) EUs; lines pass through the origin with slopes of the median values of the Ag:In ratio.

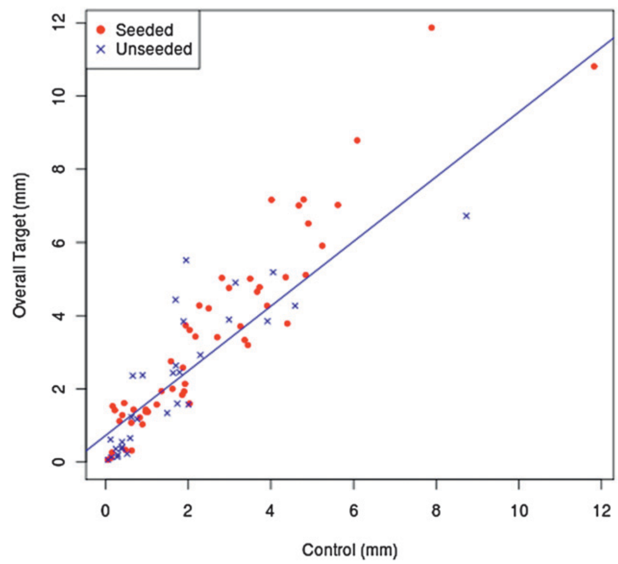

FIG. 4. Scatterplot of precipitation in overall target and control areas for seeded (red) and unseeded (blue) EUs with more than 45 generator hours; the regression line for unseeded EUs is shown in blue.

assumed that any seeding impact is directly related to the amount of seeding material, we can infer that EUs with more generator hours should have a greater seeding impact (in the absence of other confounding effects). We first note that a KS test of the distribution of generator hours in seeded and unseeded EUs shows that they are similar; that is, there is no bias in the number of generator hours between seeded and unseeded cases.

We note that there are 11 EUs with less than 35 generator hours, with 9 of these seeded and 8 being suspended EUs; therefore, about 10% of the seeded EUs received less than half the maximum possible dose of seeding material. In practice, low generator hours in unseeded EUs would not affect the overall seeding impact.

As an initial test of the sensitivity of the seeding impact to the number of generator hours, the primary precipitation analysis is repeated for the overall target area using only EUs with more than 45 generator hours. Figure 4 shows a scatterplot of the precipitation in the overall target and control area for each EU with more than 45 generator hours. In this case, the number of valid EUs drops from 107 to 84 with 53 seeded and 31 unseeded EUs. Thus, the seeding ratio is 1.6 (rather than 2) and the correlation between the overall target and control areas is found to be 0.91 (rather than 0.90); the sensitivity of the analysis is therefore expected to be a little higher than for the initial analysis using all EUs.

The analysis of seeding impact in the overall target area gives a TFI of 14% at the 3% significance level; that is, the seeding impact is higher than for the initial analysis (9%) and the significance level is much greater (13%). This result, especially when combined with the primary analysis, confirms that there is a positive impact of seeding, which is related in some way to the amount of seeding material dispersed. When repeating this for the primary target area, the TFI is found to be 14% at the 8% significance level. That is, there is a consistent increase in precipitation of 14% in both the primary and overall target areas. Bearing in mind the uncertainties in all the statistical results, the consistency of the estimated increase in precipitation may suggest that the seeding material has been effectively dispersed over both the primary and secondary target areas and that the microphysical response is essentially consistent across the region.

In summary, the seeding impact is found to be very dependent upon the number of hours of operation of the seeding generators during EUs. By removing EUs with less than 45 generator hours, the TFI in precipitation in seeded EUs in the overall target area is found to be 14% with a significance of 3%. This result highlights the need to ensure that the analysis of the seeding impact is not compromised by the inclusion of EUs with only weakly seedable conditions or with poor seeding coverage.

4. Indicators of seeding impact

The primary analysis of seeding impact is focused on the average response of the precipitation in the target area to seeding. Such analyses are needed to ensure that a robust and statistically significant seeding impact is identified. However, in order to understand the physical basis of the primary results, it is necessary to investigate the effects of seeding during individual events. In these secondary analyses, we are seeking indications of links between seeding impacts and feasible physical processes, rather than just clear statistical signals. These analyses will at least assist us in gaining a better understanding of the physics of any seeding impact and at best may allow some prediction capability to estimate the potential seeding impacts from external variables, such as wind speed or supercooled liquid water (SLW).

To study the sensitivity of the seeding impacts to variations in external variables, we introduce the local precipitation residual (RS) as an indicator of seeding impacts for an individual EU, as defined in section 2a. The variable RS is the difference between the actual precipitation in the target area and the expected or natural precipitation as estimated by regression from the control precipitation. The approach is similar to that of Super (1986) in seeking links between seeding impacts and external variables. The normalized sum of these residuals is the TFI, which is the overall indicator of a seeding impact in the primary analysis.

To better understand the physical basis of the apparently positive (but weak) impacts of seeding on precipitation, we will analyze statistical relationships between the indicator RS and physical variables that are expected to have an impact on seeding efficacy. A particular investigation is of the differences between the relationships in seeded and unseeded EUs. KS tests are used to identify differences in distributions between seeded and unseeded cases, and Pearson's product moment correlations (Everitt and Hothorn 2006) are used to identify linear relationships between seeding indicators and physical variables. In these statistical tests, we will assume that a level of 10% or better is worthy of consideration for physical (as well as statistical) significance.

Another indication of seeding impact due to variations in a given variable is derived by repeating the primary precipitation analysis with upper or lower thresholds on the value of that variable. The indicator in this case is the TFI. In this way we can examine the change in TFI when, for example, only EUs with high values of SLW are included in the primary analysis.

a. Generator hours

In section 3, we found that there is a strong relationship between seeding impact and the number of generator hours for seeded EUs. This result is emphasized in Fig. 5, which shows the impact on the TFI when EUs below a specified threshold level of generator hours are ignored. The bottom panel in Fig. 5 gives the statistical

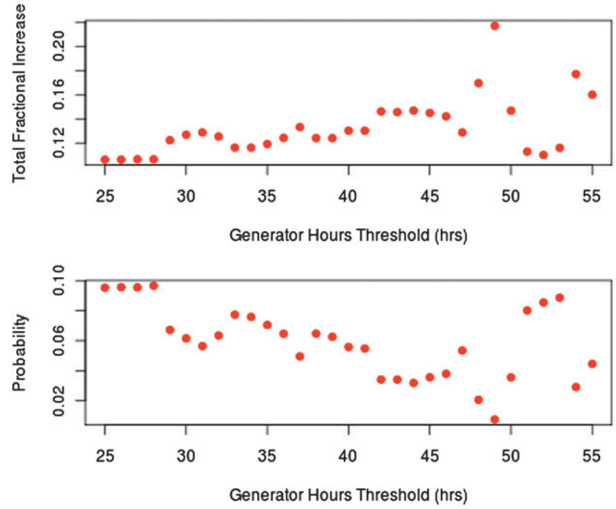

FIG. 5. Variation in TFI in seeded EUs with a lower threshold on generator hours: (top) TFI vs the threshold and (bottom) significance level of the increase.

significance of the estimate of TFI for each case. It is seen that TFI does not change very much as the threshold increases from low values, which suggests that EUs with very low generator hours make a negligible contribution to TFI. More particularly there is essentially a linear increase in TFI as the threshold is increased up to about 45 h. The statistical significance of the estimate also increases (probability decreases) as the threshold increases up to 45 h. Beyond about 45 h, the number of EUs falls rapidly and so the results become unreliable. The jumps in the values in Fig. 5 occur as the number of samples changes with the threshold. Inspection of Fig. 5 suggests that the optimal threshold value for generator hours could be around 42 h.

We now use all EUs to study the statistical relationship between generator hours and the seeding indicator. The precipitation residual (RS) is different in seeded and unseeded EUs at the 6% level. Moreover, while unseeded residuals are uncorrelated with generator hours, the residuals in seeded EUs have a correlation of 0.20 with generator hours at the 9% significance level.

In summary, there is no response of precipitation in the target area to the number of generator hours in unseeded EUs, but there is a consistent increase in precipitation for seeded EUs as the number of generator hours is increased. In particular, the TFI is found to increase steadily as a lower threshold of up to 45 h is used in the analysis. The correlation between target precipitation residual (RS) and generator hours is 0.20 at the 9% significance level. On the basis of the relationship between seeding impact and generator hours, and in order to maximize any signal relating seeding impact to other physical variables, only the 84 EUs with more than 45 generator hours are used in the following analyses.

b. Supercooled liquid water

The seeding hypothesis for SPERP is based on the assumption that the seeding material is able to transform SLW into ice, leading to enhanced precipitation at the ground. As we have found a positive relationship between the number of generator hours and seeding impact at both the bulk and event levels, it is expected that there should be some relationship between SLW and local seeding impact.

There are two key SLW variables analyzed from the radiometer at Blue Cow. SLWS is the mean observed SLW over the 0.5 h before an EU commences. Indeed, SLWS needs to be 0.05 mm or greater (or ice trips need to have been observed at the surface) to satisfy a basic criterion for seeding onset. SLWM is the mean value of SLW over the EU evaluation period.

As expected, there is no difference between the distributions of SLWS in seeded and unseeded EUs. For the precipitation residuals, there is a correlation of 0.35 at the 1% level with SLWS in seeded EUs and none in unseeded EUs. Thus, there is a clear relationship between SLW at the start of an EU and the local seeding impact. It is interesting to note that, although it is not statistically significant, there is a weak negative correlation (-0.13) between the precipitation residual and SLWS for unseeded EUs; that is, in the natural cases, high values of SLWS tend to lead to precipitation that is less than the expected value when other factors are fixed. These results provide support for the seeding hypothesis that the seeding material leads to the consumption of SLW to produce precipitation.

The importance of the level of SLWS is further demonstrated by Fig. 6, which shows the TFI when an upper threshold is placed on SLWS; that is, as the threshold is reduced, EUs with SLWS greater than the threshold are ignored in the analysis. It is clear that the TFI essentially increases linearly with SLWS. This result suggests that the EUs with high SLWS have the greatest impact. Further support for this conclusion is given by a similar analysis but with a lower threshold on SLWS; it is found that there is little change in TFI as the threshold is raised to 0.2 mm; that is, the low-SLWS EUs are not contributing substantially to the TFI. We note that the statistical significance of these analyses may not be high (as indicated by the noisiness of the plots), but they do provide support for the physical basis of the key results. This result is not inconsistent with the aircraft observations of Ryan and King (1997), who find that the rainfall enhancement tends to increase as SLW increases from about 0.2 to 0.5 g m^{-3} over a 5-min-averaging period; for 1 km of uniform cloud, 0.2 g m^{-3} corresponds to 0.2 mm of SLW.

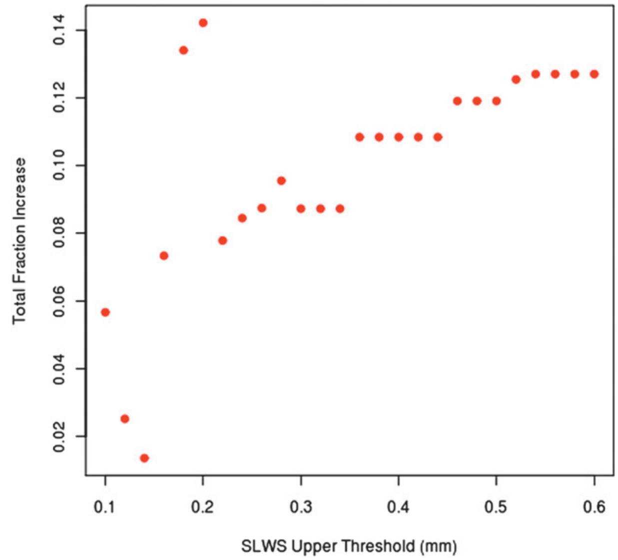

FIG. 6. Variation in TFI with an upper threshold on SLW in the 0.5 h before the start of an EU; i.e., EUs with SLWS greater than the threshold are ignored in the analysis of the seeding impact.

The physical relationship between SLW and precipitation is further emphasized by the study of SLWM, that is, the mean SLW over the duration of an EU. We first note that, although a KS test shows that the seeded and unseeded distributions of SLWM are not significantly different, it is apparent that the seeded values (median of 0.063 mm) tend to be lower than the unseeded values (median of 0.091 mm). There is a correlation of 0.64 at the 0.00% level between RS and seeded SLWM, and none for unseeded EUs.

Although there is a positive correlation between the residuals and SLWM, the relationship is quite complex. By repeating the analysis of TFI with a lower threshold on SLWM, we find that TFI is essentially constant up to about 0.1 mm but that it decreases with a SLWM threshold above 0.1 mm. This result suggests that most of the seeding impact is arising from EUs with values of SLWM less than about 0.1 mm. Indeed, if an upper threshold on SLWM is applied to the analysis of TFI, we find that the TFI tends to increase as the threshold is decreased. These results suggest that the substantial effect of SLWM on precipitation is associated with low values of SLWM. This is consistent with the initial observation that seeded EUs tend to have lower values of SLWM than do unseeded EUs.

In summary, there is a strong positive correlation between the precipitation residuals in seeded EUs and SLW both at the start of an EU (associated with high values) and SLW averaged over the whole EU (associated with low values). These results suggest that seeding leads to the consumption of SLW.

The relationship between seeding impact and SLW over an EU is explored further by consideration of DSLW, which is the difference between SLW at the start of an EU (SLWS) and the mean value of SLW over an EU (SLWM). From the analyses of SLWS and SLWM individually, we would expect a positive seeding impact to be indicated by a positive value of DSLW; that is, seeding should lead to the consumption of SLW. Indeed, when repeating the analysis of TFI with a lower threshold on DSLW, it is found that TFI is essentially constant for negative values of DSLW, but that it increases steadily for positive values of DSLW.

Further analysis of the relationship between precipitation residuals and DSLW for individual EUs shows that for DSLW < 0 the relationship is the same as for unseeded EUs (i.e., a strong negative correlation). For positive DSLW, the correlation is lost for seeded EUs. It would seem that the observation of a positive value of DSLW for an EU strongly suggests a positive seeding impact.

In summary, for individual EUs and across all EUs, we find moderately strong relationships between SLW and the seeding impact. The higher the value of SLW at the start of an EU, the greater is the seeding impact. Moreover, there is clear evidence of seeding impacts being associated with the consumption of SLW over an EU.

c. *Wind*

The wind at the $-5°$C level is taken as the variable of interest: WSP5C denotes the wind speed (m s^{-1}) and WDR5C denotes the wind direction (°). The KS test shows that both wind speed and direction have similar distributions for seeded and unseeded EUs.

For wind direction, there is a positive correlation between RS in seeded EUs and WDR5C of 0.32 at 2% significance, while there is no correlation for unseeded EUs. Repeating the seeding impact analysis of TFI with a lower threshold on WDR5C, it is found that there is little impact of the threshold on TFI for WDR5Cs less than about 270°, but there is a steady increase in TFI beyond 270°. This result suggests that the northerly component of the wind is important to maximize the seeding impacts. The main range in the target area tends to have a northeasterly alignment, and so a reasonable a priori assumption is that the sensitivity to wind direction is related to the optimization of orographic lifting. Indeed, Warburton and Wetzel (1992) report that peak values of SLW at Blue Cow occur in northwesterly flows. It is difficult to make direct comparisons between these results for the Snowy Mountains and those of Long and Huggins (1992) at the Baw Baw Plateau to the south, because the orientation of the local ranges greatly affects the optimal flow for the enhancement of SLW.

For both seeded and unseeded EUs, there is a positive correlation between RS and wind speed. This result suggests that high wind speeds are associated with higher-precipitation events for both natural and seeded conditions. The correlation of seeded residuals and WSP5C is 0.48 at 0.03% significance, which is higher than with unseeded EUs where the correlation is 0.32 at the 8% level. We note that Super (1986) finds a correlation of 0.17 between the wind speed normal to the barrier range and the target precipitation in unseeded 6-h events.

Repeating the seeding impact analysis for TFI with a lower threshold on WSP5C, it is found that up to about 8 m s^{-1} there is little impact on the TFI. Between about 8 and 18 m s^{-1} there is a steady increase in TFI and then it drops off at higher wind speeds. Thus, there appears to be an optimum range of wind speeds for obtaining a positive seeding effect. A wind speed of 20 m s^{-1} allows about 15 min for material to travel from the generators to Blue Cow, which is at the lower limit of travel times suggested by Dennis (1980) for ground-based seeding. There is a very significant correlation of 0.35 between wind speed and SLWS, supporting the concept that the generation of SLW in orographic conditions is enhanced by higher wind speeds. This correlation suggests that the weak seeding impact at low wind speeds is associated with lower values of SLW.

d. *Cloud-top temperature*

It has been observed that natural ice-precipitation processes tend to occur at temperatures below about $-20°$C, and so seeding is expected to be especially effective for cloud-top temperatures up to that value (Dennis 1980). The SPERP criterion for seeding onset is that the cloud-top temperature (CTT) must be less than $-7°$C, because the seeding material is most effective at lower temperatures.

The statistical analysis of CTT shows that the distributions in seeded and unseeded EUs are very similar. There is no correlation between RS and CTT in seeded EUs, while they are strongly negatively correlated in unseeded EUs with a correlation of -0.37 at the 4% level. Super (1986) also reports a correlation of -0.38 between CTT and the target precipitation for unseeded conditions. This suggests that, in the absence of seeding, precipitation tends to increase as CTT decreases; that is as clouds deepen. However, this effect is masked in seeded clouds, because large residuals are unlikely for very low CTTs.

Repeating the seeding impact analysis of TFI with a lower threshold on CTT, it is found that the TFI has a slight upward trend as the CTT threshold is increased to about −20°C, where it reaches a maximum of about 20%. For higher thresholds, the TFI steadily decreases to negative values. These results suggest that the CTT should be as low as possible up to −20°C, after which there is little or even a negative impact. Although there is no correlation between CTT and SLW at the start of an EU, the positive impact up to −20°C is likely to be due to the increasing depth of the SLW. High values of SLW do not occur for very cold clouds. The result that seeding appears to have some impact at CTTs as low as −20°C is consistent with an analysis of seeding impacts over the Bridger Range in Montana (Super 1986), where events of 6-h duration were considered. On the other hand, Ryan and King (1997) report that there are few seeding opportunities in Tasmania for CTTs lower than −15°C; the reason for this difference is not clear.

e. Cloud heights

Statistical analyses can be carried out on the relationship between RS and the height of the −5°C level (HGT5C), the freezing level (HGT0C), cloud-top height (HGTCT), and cloud-base height (HGTCB) at the start of each EU. We first note that HGTCT is expected to have similar relationships as CTT, and it is found that there is a negative correlation between HGTCT and RS for unseeded EUs (−0.30 at 9%). Second, cloud base often reaches the ground and so statistical relationships are hard to identify for HGTCB. Third, since no significant correlations are found between the freezing level and RS, we concentrate on HGT5C.

HGT5C is expected to be relevant because silver iodide tends to be effective below −5°C (Dennis 1980). A KS test shows that the distributions of HGT5C are similar for seeded and unseeded EUs, and so other statistical differences should be due to seeding effects. For unseeded EUs, there is not a significant correlation between HGT5C and RS. However, there is a correlation of 0.26 at the 6% level between HGT5C and RS for seeded EUs.

This correlation between HGT5C and RS for seeded EUs is confirmed by repeating the seeding impact analysis for TFI with an upper threshold on HGT5C. It is found that the TFI steadily decreases as the threshold is lowered. Similarly, a lower threshold on HGT5C in the seeding impact analysis suggests that most of the impact arises from EUs with HGT5C above about 1800 m, which is around the peaks of the mountain ranges. The physical basis of this result may be that the higher HGT5C is, the greater the depth of cloud available for ice (formed from SLW above HGT5C) to sweep out cloud particles below HGT5C. Further research is needed to clarify this apparent relationship.

A criterion for seeding onset is that the difference between cloud top and the −5°C level must be greater than 400 m, and so it is appropriate to consider the sensitivity of seeding impacts on this height difference (DHCT5C). In fact, there is a correlation of −0.98 between DHCT5C and CTT. This correlation is even higher than the correlation between HGTCT and CTT, because HGT5C also has a small (−0.17 at 9% significance) correlation with CTT; that is, the height of the −5°C level tends to rise with cloud-top height for the systems suitable for seeding. The strong correlation between DHCT5C and CTT suggests that the relationship between DHCT5C and RS should be very similar to that with CTT. Indeed, we find that there is a correlation of 0.35 (6% significance) between RS for unseeded EUs and DHCT5C, but no significant correlation for seeded EUs. Analysis of the TFI with upper and lower thresholds on DHCT5C suggests that the main contributions to the seeding impact come from events where DHCT5C is less than about 2500 m.

f. Precipitation

It is useful to explore the relationship between seeding impact and the absolute level of precipitation. There is no significant correlation between RS and the precipitation in the control area (PRCON) for unseeded EUs, while for seeded EUs the correlation is 0.47 at the 0.04% level. This result suggests that the seeding impact increases with the level of precipitation, and it is confirmed by an analysis of the total seeding impact with a lower threshold on PRCON. As the threshold increases from 0, the TFI increases steadily, such that the TFI is about 21% for a cutoff of 2 mm. It is apparent that the seeding impact increases with increasing precipitation of seedable events in the control area, which have a median precipitation rate of only about 0.4 mm h^{-1}. It is noted that the heaviest-precipitation events will generally not be seedable; the maximum 5-h precipitation over the primary target for the winters of SPERP corresponded to 6 mm h^{-1}. The strong correlation between RS and PRCON suggests that the seeding impact is largely multiplicative, as is essentially assumed when a double-ratio analysis is carried out (Smith et al. 1963). On the other hand, Reynolds (1988) summarizes direct measurements of seeding impacts to suggest that seeding leads to additive increases of order 0.1–0.6 mm h^{-1}. Averaged over the whole target area and over the 5 h of an EU, the 14% estimated impact of SPERP (section 3) on a natural precipitation rate of 0.4 mm h^{-1} yields an enhancement rate of 0.06 mm h^{-1}, which is at the lower limit of the Reynolds results. The results of this technical analysis of SPERP are providing the basis for a separate investigation of the economic feasibility of cloud seeding in the Snowy Mountains.

TABLE 1. Correlations between independent variables and seeding impact, represented by precipitation residual (RS).

Variable	Correlation with precipitation residual for seeded EUs (significance level)	Comment
Generator hours	0.20 (9%)	Seeding impact increases with amount of seeding material released
SLW at start of EU (SLWS)	0.35 (1%)	Seeding impact increases with available SLW
Mean SLW over EU (SLWM)	0.64 (0.00%)	There are few seeded EUs with SLWM greater than 0.1 mm; i.e., seeding impact is associated with low values of SLWM
SLWS − SLWM (DSLW)	−0.71 (0.2%) for DSLW < 0	For SLWS < SLWM, seeded EUs behave like unseeded EUs; for SLWS > SLWM, there is no correlation for seeded EUs, suggesting that SLW is consumed when seeding is effective
Wind direction at −5°C level	0.32 (2%)	Seeding is most effective for northwesterly winds, perhaps owing to the optimization of the alignment of the wind with the main range in the target area
Wind speed at −5°C level	0.48 (0.03%)	The correlation seems to be associated with a window between about 8 and 18 m s^{-1}, such that winds too low do not generate substantial amounts of SLW and winds too high do not allow sufficient time for microphysical effects to occur
Cloud-top temperature (CTT)	No significant correlation	For unseeded EUs, there is a correlation of −0.37 at the 4% level between CTT and the precipitation residual; this is lost for seeded EUs, when high residuals do not occur at very low CTTs
Cloud-top height (HGTCT)	No significant correlation	For unseeded EUs, there is a correlation of 0.31 at the 9% level between HGTCT and precipitation residual; this is lost for seeded EUs
Height of −5°C level (HGT5C)	0.26 (6%)	The higher HGT5C is, the deeper the layer for ice particles to sweep up cloud particles below
Diff between cloud-top height and −5°C level (DHCT5C)	No significant correlation	For unseeded EUs, there is a correlation of 0.35 at the 6% level between DHCT5C and the precipitation residual; this is lost for seeded EUs
Precipitation in control area	0.47 (0.04%)	Seeding effect is essentially multiplicative

g. *Summary*

It has been found that there are significant correlations between the indicators of a seeding impact and variables that are expected to affect the efficacy of seeding. Table 1 summarizes the relationship between the local precipitation residual and each variable. The correlations in the table are generally not large. However, the important feature is that all the listed correlations are significant at the 10% level at least; that is, these correlations are very unlikely to be due to chance. It is therefore appropriate to use the correlations for physically based analyses of a seeding impact, as suggested by Super (1986).

The results summarized in Table 1 are consistent with the seeding hypothesis that silver iodide leads to enhanced precipitation in the target area through the transformation of SLW into ice. The microphysical hypothesis is supported especially by the correlations of the seeded precipitation residuals with generator hours, SLW, and height of the −5°C level. The relationship of the seeding impact with cloud-top temperature is also consistent with the seeding hypothesis, because the natural negative correlation is masked. The relationship of the precipitation residual with the wind speed suggests that there is an optimal window between about 8 and 18 m s^{-1}. The results of SPERP are consistent with those of Super (1986) from the Bridger Range Experiment to the extent that the seeding impact is found to be sensitive to external variables, such as cloud-top temperature, wind speed, and wind direction.

5. Conclusions

A 5-yr cloud seeding project in the Snowy Mountains of southeastern Australia has been completed under carefully designed and controlled conditions. The primary analysis leads to a positive, but not statistically significant, impact of seeding on precipitation; in particular, there is a 7% increase in precipitation across the primary target area but with only a 24% significance level. On the other hand, the primary analysis of the targeting of seeding material is unequivocally successful, with the maximum level of Ag in the primary target area found to be significantly greater in seeded than unseeded EUs. A secondary analysis of the ratio of Ag:In supports the proposal by Chai et al. (1993) that there are enhanced levels of seeding material compared with a passive tracer in seeded EUs owing to AgI acting to nucleate additional ice particles. This result implies that there has been a microphysical impact on cloud precipitation processes, although that impact was not confirmed in the primary precipitation analysis.

Another secondary analysis shows that a substantial source of uncertainty in the estimation of the impact of seeding on precipitation is associated with EUs where the seeding generators operated for a relatively small number of hours. Indeed, 10% of EUs had fewer than half the maximum number of generator hours. Most importantly, when the analysis of the seeding impact is repeated using only EUs that received more than 45 generator hours, the increase in precipitation in the primary target area is 14% at a significance level of 8%, which is within the specified level of 10%. When that analysis is applied to the overall target area, the precipitation increase is 14% at the 3% significance level.

Further, secondary analyses reveal that seeding had an impact on virtually all of the physical variables examined. In particular, the seeding impact is influenced by the amount of SLW available at the start of an EU, supporting the concept that SLW tends to be consumed in seeded (but not unseeded) EUs. The seeding impact is also affected by wind speed and direction, as well as the height of the $-5°C$ level. All of these relationships are statistically significant and are consistent with the seeding hypothesis. The relationships are strengthened by the observation that the randomization strategy generally led to similar distributions of the independent variables in seeded and unseeded EUs.

Acknowledgments. This evaluation of SPERP is based on datasets that were carefully collected and preprocessed by Snowy Hydro Limited officers. The thoughtful comments of the project referees, Dr. Warren King and Dr. Doug Shaw, helped to refine the analysis. The manuscript benefited greatly from the insights of three anonymous reviewers.

REFERENCES

Bauer, D. F., 1972: Constructing confidence sets using rank statistics. *J. Amer. Stat. Assoc.,* **67,** 687–690.

Chai, S. K., W. G. Finnegan, and R. L. Pitter, 1993: An interpretation of the mechanism of ice-crystal formation operative in the Lake Almanor cloud-seeding program. *J. Appl. Meteor.,* **32,** 1726–1732.

Dennis, A. S., 1980: *Weather Modification by Cloud Seeding*. Academic, 275 pp.

Everitt, B. S., and T. Hothorn, 2006: *A Handbook of Statistical Analyses Using R*. Chapman and Hall, 269 pp.

Huggins, A. W., S. L. Kenyon, L. Warren, A. D. Peace, S. P. Bilish, and M. J. Manton, 2008: The Snowy Precipitation Enhancement Research Project: A description and preliminary results. *J. Wea. Modif.,* **40,** 28–51.

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.

Manton, M. J., J. Dai, and L. Warren, 2009: Evaluation plan for a snow enhancement experiment in Australia. *J. Wea. Modif.,* **41,** 59–74.

——, L. Warren, S. L. Kenyon, A. D. Peace, S. P. Bilish, and K. Kemsley, 2011: A confirmatory snowfall enhancement project in the Snowy Mountains of Australia. Part I: Project design and response variables. *J. Appl. Meteor. Climatol.,* **50,** 1432–1447.

Marsaglia, G., W. Tsang, and J. Wang, 2003: Evaluating Kolmogorov's distribution. *J. Stat. Softw.,* **8,** 18.

Nicholls, N., 1989: Sea surface temperatures and Australian winter rainfall. *J. Climate,* **2,** 965–973.

Reynolds, D. W., 1988: A report on winter snowpack-augmentation. *Bull. Amer. Meteor. Soc.,* **69,** 1290–1300.

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.

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

Super, A. B., 1986: Further exploratory analysis of the Bridger Range winter cloud seeding experiment. *J. Climate Appl. Meteor.,* **25,** 1926–1933.

Warburton, J. A., and M. A. Wetzel, 1992: Field study of the potential for winter precipitation enhancement in the Australian Snowy Mountains. *Atmos. Res.,* **28,** 327–363.