

AN ANALYSIS OF CLOUD SEEDING OPERATIONS OVER TASMANIA FROM 1964 TO 2005

Anthony Morrison¹, Steven Siems¹, Michael Manton¹, Alex Nazarov²

¹School of Mathematical Sciences, Monash University, Melbourne, Australia

²Hydro Tasmania, Hobart, Australia

1. INTRODUCTION

One of the common, long-lasting criticisms of cloud seeding is that there is little convincing evidence that this practice increases precipitation at the surface. This failure does not so much lie with the theory of cloud seeding, but with the challenge of observing rainfall. The high natural variability of rainfall across various spatial and temporal scales means that it is impossible to accurately forecast it, and if one cannot predict with any certainty the amount of rainfall from a given system, then it is not possible to quantify any change to that system. This problem has repeatedly been attacked by long-term statistical analysis with limited success. Even over an extended period of time (months to years), the natural variability of rainfall remains too great to definitively quantify any changes to precipitation. Moving to any longer time scale leads to pragmatic operational issues.

The central plateau of Tasmania (Figure 1) is relatively unique in that this wilderness region has served as the target area of three independent cloud seeding experiments, as well as two operational cloud seeding programs. In all between the years of 1964 to 2005, a total of 23 of the 42 winter seasons were seeded either in an experimental mode or operationally. As this region has seen no significant development over this time period, it is in many ways an ideal location to study the effects of cloud seeding over an extended period of time.

The first cloud seeding research activity (intermittent from 1964-1971) was reported to have produced rainfall increases of up to 30% (Ryan and King, 1998) in the autumn. In this period silver iodide was applied 30 to 60 minutes upwind of the target region. A randomization strategy was implemented in which pairs of 12-day blocks were defined with one of the 12-day blocks being seeded and the other not. In addition to this, the early

part of this field work alternated years for seeding. Having gained insight from this initial seeding operation, a second seeding experiment was conducted from 1979-1983. During this particular period, randomization occurred at the individual “suitable seeding day” level with an overall seed/no seed ratio of 2:1. The CSIRO reported back to the Tasmanian Hydro-Electric Authority (Shaw 1984) that a rainfall increase of 37% was found on suitable seeding days. The Tasmanian Hydro-Electric Authority embarked on an operational seeding exercise from 1988 to 1991 in response to relatively low reservoir levels. No control site was established during this period, and no effort to quantify the effect of this seeding was undertaken. From 1992 to 1994 the CSIRO conducted a third experimental seeding effort over Tasmania in which dry ice was used as a seeding agent instead of silver iodide. In other respects, the experiment was conducted under the guidelines of the second research period. No results from this research have been published. Finally, beginning in 1998, an operational cloud seeding program has been maintained by the Hydro Tasmania (the successor to the Hydro-Electric Corporation of Tasmania.)



FIGURE 1. Map of catchment area over Tasmania.

Taken individually, these different seeding efforts have suffered from the standard uncertainty arising from the high natural spatial and temporal variability of rainfall. Moreover, it is difficult to define an appropriate control area from which to quantify the effect of cloud seeding. The control areas defined for the second and third research phases differ from the original control areas. Pook & Budd (2002) conducted a review of these efforts for Hydro Tasmania and concluded that “there is concern that the significance of these apparent increases might not be sufficiently robust because of the natural variability in the ratios of target to control rainfall and the small sizes of the samples.” Taken collectively, however, these five seeding periods span a period of 42 years with a total of 19 years that were not seeded at any time. While the cloud seeding activities in Tasmania have not been designed or implemented as a single long-term experiment, it is possible to analyse them as such.

2. ANALYSIS STRATEGY

The Bureau of Meteorology maintains a record of over 1000 stations across Tasmania dating back to the late 1800's. These individual records come and go as sites are moved or equipment upgraded, but as one might expect, the number of sites operating at any given time has steadily increased over time. In addition to these sites, Hydro Tasmania maintains approximately 90 additional sites largely up in the catchment area and immediate surroundings.

As the purpose of any useful cloud seeding project is to increase the annual or at least seasonal rainfall to a measurable extent over an area of economic significance, we take a calendar month as the basic unit of time. The simplicity in such an approach offers a number of advantages; most notably there is no ambiguity in our definition of seeding periods. If the aircraft made even one single seeding mission in a given month, we have defined that month as a seeded month. The records show that, with the possible exception of early autumn months, if the seeding aircraft was available, it operated at least once in that month. The drawback to such a conservative definition is an obvious weakening of any signal. If the aircraft only operated for a short period in a given month, it stands to reason that it would not have a significant impact on the overall monthly rainfall.

Turning to the precipitation, average monthly rainfall totals are calculated over the entire target area. This is the region of economic significance. Unfortunately, this adds another potential complication to our analysis; the target region has been increased from the original research operation in 1964 to the operational target areas of 2005 (Figure 1). By choosing the larger target area we have again potentially weakened any signal that seeding may have on the rainfall observations. In practice, however, the effect of this expanded target region is thought to be limited since the vast majority of the long-lived surface sites were placed within the original target area.

The rainfall over the target area is defined by simply taking an arithmetic average of the available data for a given month. This approach has two difficulties that we need to consider. The first is that the spatial distribution of rain-gauge sites may not be adequate to properly sample the spatial variability, and so our estimated rainfall may not be accurate. The second is that the number of operating rain gauge sites varies with time. Indeed, there is not a single rain gauge that has operated continuously in the target region from 1964 through 2005. This intermittency in observations can readily lead to a hidden sampling bias in the analysis. If, for example, a site operated at a high elevation it would most likely be wetter than a site at low altitude. If this site were to have operated for a limited segment of the overall period, then it would arbitrarily enhance the monthly area averaged precipitation for the years when it was in operation, whether seeded or not.

We explore the sensitivity to this in two ways. First we objectively exclude rain gauges based on the length of time they have been in operation. We create a range of monthly, area-averaged precipitation tables based on the rain gauges that have operated for a minimum of at least 75% (most restrictive), 50% and 25% (least restrictive) of the period of analysis. Second, we calculate a sampling error directly, which is discussed shortly.

The sensitivity of the definition of the monthly area-averaged rainfall is explored further by considering only those sites maintained by the Bureau of Meteorology (BoM) versus records that include both the BoM sites and the Hydro Tasmania (HT) sites. Finally, for the upcoming Monte Carlo experiment, we also consider extending the overall analysis period from

1950-2005, as opposed to a more restrictive 1964-2005. This, ultimately, left us with 12 variations for the definition of the monthly area-averaged target rainfall depending on the inclusion or exclusion of sites within the target region. It was comforting to find that the basic results of our analysis were not particularly sensitive to these different definitions.

3. MONTE CARLO ANALYSIS

Before the analysis begins, it is compelling to visualise the difference between the seeded and unseeded months (Figure 2.) It is perhaps startling to see such a strong difference given the conservative nature of our definition of area-averaged rainfall. This difference is strongly observed over the winter months (JJA) and the spring (SON), but not really evident in the summer or autumn. We recall that the early analysis found the strongest results during the autumn months.

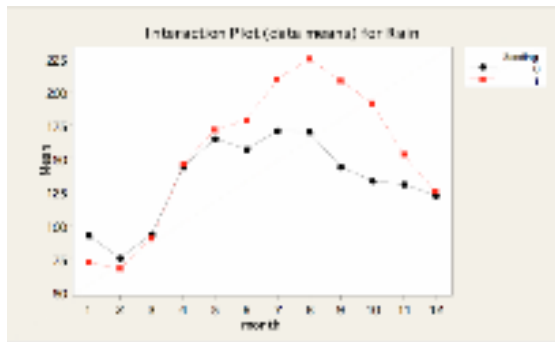


FIGURE 2. Time series of monthly rainfall when averaged for seeded and unseeded periods. Use

50% threshold for ALL target sites from 1964-2005.

As a first step, an analysis of variance (ANOVA) test is undertaken to see if the overall variability may be attributed to various factors. Here we explore the variance associated with the year, the month and whether or not seeding was undertaken. For the 50% threshold for all target sites (1964-2005), there is less than an 0.001% change that this variability would be present if seeding were of no impact. The month was also observed to be significant in explaining the observed variability. The year, too, was found to be significant, but not at the extreme level as the month or the presence of seeding.

A more crude but direct test of the importance of the seeding is to assume that it is of no consequence to the observed rainfall. If so then the application of a random seeding mask should have little impact measured “seeded” precipitation. Indeed, a Monte Carlo experiment may be run on testing whether the seeding mask has any impact on the defined seeding total. We conduct this randomisation a total of 10,000 times. If seeding were of no importance than the true seeding realisation would most likely lie within one standard deviation of the mean of this Gaussian distribution. In this randomisation we ensure that the total number of months seeded over the period from 1950-2005 remains conserved. As the real seeding mask contains 23 August seedings, each of the 10,000 randomised masks will have 23 August seedings.

Cutoff	Period	Sites selected	No sites	True Precip	Avg. MC Precip	MC StDev	StDev increase	Percent increase
50%	1964-2005	BoM only	16	1661	1619	57.6	0.73	2.6
50%	1964-2005	BoM and HT	32	1711	1653	57.0	1.03	3.6
75%	1964-2005	BoM only	5	1734	1770	64.9	-0.55	-2.0
75%	1964-2005	BoM and HT	19	1782	1730	59.9	0.87	3.0

TABLE 1. Monte Carlo Analysis for full year

Cutoff	Period	Sites selected	No sites	True Precip	Avg. MC Precip	MC StDev	StDev increase	Percent increase
50%	1964-2005	BoM only	16	546	504	16.2	2.53	8.2
50%	1964-2005	BoM and HT	32	568	527	15.8	2.54	7.6
75%	1964-2005	BoM only	5	570	534	18.4	1.96	6.8
75%	1964-2005	BoM and HT	19	5.85	549	16.8	2.16	6.6

TABLE 2. Monte Carlo Analysis for Winter only

Cutoff	Period	Sites selected	No sites	True Precip	Avg. MC Precip	MC StDev	StDev increase	Percent increase
50%	1964-2005	BoM only	16	516	447	19.2	3.55	15.2
50%	1964-2005	BoM and HT	32	515	450	18.6	3.48	14.4
75%	1964-2005	BoM only	5	569	503	22.1	2.96	13.0
75%	1964-2005	BoM and HT	19	543	475	19.9	3.42	14.4

TABLE 3. Monte Carlo Analysis for Spring only

Taken at face value, these results strongly suggest that cloud seeding is leaving a signature on the precipitation over the target region during the spring and winter. A three standard deviation increase suggests that only about 10 of the 10000 trials resulted in more precipitation than the true seeding mask. A further analysis of control regions outside of the target area suggests that while this signature is quite impressive, some signature is evident over a far wider region. It is unlikely that this result can be taken strictly at face value.

We briefly return to the potential problem of sampling error, i.e., the error that may potentially arise when sites are used that only run for a limited part of the time period. For this test we consider each individual rain gauge site individually. For each month, for each site, we calculate the mean and standard deviation. We then return random precipitation amounts back to these sites based on the mean and standard deviation. We are, in essence, returning observations that are completely unbiased to whether seeding had occurred or not. Thus if our results were influenced by the inclusion of particularly “wet” sites during periods of cloud seeing, this should provide a sampling bias evident in this idealised, unbiased data as well as the original results. We create 100 such idealised, unbiased data sets for each station/month. Feeding this into the Monte Carlo analysis leaves us with a means to define a sampling bias. The sampling bias is most strong and positive when the time period runs from 1950-2005. This is actually because on a few BoM sites are present in the target area over during the 1950s. The inclusion of these data actually return stronger results for the Monte Carlo test, but this stronger signal is matched in the sampling bias found. Hence, we have focussed on the time period of 1964-2005. The second conclusion reached from the sampling bias analysis is that the sampling bias is smallest at the highest (most selective) thresholds. When limiting ourselves to the cases presented in Table 3, we

note that the sampling bias is less than 0.5 standard deviations and cannot account for the extreme signal detected.

4. CONTROL ANALYSIS

One final consideration must be explored, especially as the seeding years are not truly randomised. Are these results in response to some kind of regional climatological patterns? If the overall precipitation were increasing in time, the Monte Carlo analysis would not be able to account for this. Consider the extreme precipitation changes observed in the Sahel in Africa or in Western Australia. If the reverse of this were present over Tasmania, it would increase precipitation over the later years and bias the MC results.

To test for this, we undertake this Monte Carlo analysis to a number of control regions within Tasmania, but not connected to the target area. These results suggest that the eastern part of Tasmania experiences approximate a one standard deviation increase in precipitation while a control area defined to the west of the target exhibits over two standard deviations (increase) in precipitation. The potential implications of such findings will be more fully discussed.

REFERENCES

- Pook, M. J., and W. F. Budd, 2002: An evaluation of cloud seeding operations carried out in Tasmania by Hydro Tasmania. CRC for Antarctica and the Southern Ocean, Hobart.
- Ryan, B. F. and W. D. King, 1997: A critical review of the Australian experience in cloud seeding. Bull. Amer. Met. Soc., 78, 239-254.
- Shaw, D. E., W. D. King and E. Turton, 1984: Analysis of Hydro-Electric Commission cloud seeding in Tasmania 1979-1983. CSIRO report to the HEC, CP394, Hobart.