



UNIVERSITY
OF WOLLONGONG
AUSTRALIA

University of Wollongong
Research Online

Centre for Statistical & Survey Methodology
Working Paper Series

Faculty of Engineering and Information Sciences

2010

Accounting for Spatiotemporal Variation of Rainfall Measurements when Evaluating Ground-Based Methods of Weather Modification

Stephen Beare

University of Wollongong, uow@beare.edu.au

Ray Chambers

University of Wollongong, ray@uow.edu.au

Scott Peake

University of Wollongong

Recommended Citation

Beare, Stephen; Chambers, Ray; and Peake, Scott, Accounting for Spatiotemporal Variation of Rainfall Measurements when Evaluating Ground-Based Methods of Weather Modification, Centre for Statistical and Survey Methodology, University of Wollongong, Working Paper 17-10, 2010, 27p.
<http://ro.uow.edu.au/cssmwp/67>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library:
research-pubs@uow.edu.au



Centre for Statistical and Survey Methodology

The University of Wollongong

Working Paper

17-10

Accounting for Spatiotemporal Variation of Rainfall Measurements
when Evaluating Ground-Based Methods of Weather Modification

Stephen Beare, Ray Chambers, Scott Peak, Jennifer M Ring

*Copyright © 2008 by the Centre for Statistical & Survey Methodology, UOW. Work in progress,
no part of this paper may be reproduced without permission from the Centre.*

Centre for Statistical & Survey Methodology, University of Wollongong, Wollongong NSW
2522. Phone +61 2 4221 5435, Fax +61 2 4221 4845. Email: anica@uow.edu.au

Accounting for Spatiotemporal Variation of Rainfall Measurements when Evaluating Ground-Based Methods of Weather Modification

Stephen Beare¹, Ray Chambers², Scott Peak³, Jennifer M Ring³

¹Analytecon, 4 Johnston Street, Narrabundah, ACT 2604, Australia

²Centre for Statistical and Survey Methodology, University of Wollongong, Wollongong, Australia

³Australian Rain Technologies, Pier 8/9, 232 Hickson Road, Millers Point, NSW 2000 Australia

ABSTRACT

Weather modification trials tend to rely on randomized experimental designs. Unfortunately, these designs have so far not demonstrated sufficient power to detect a small weather modification signal against the large level of background variation in rainfall. Further, randomized experimental designs are generally not possible when dealing with ground-based sources of weather modification such as industrial pollution. Statistical modeling of rainfall gauge measurements that attempt to control for meteorological and orographic variation in rainfall measurements seem better suited in this regard. Evaluation would be relatively simple if we could separate the sources of variation into changes in meteorological conditions in time and fixed effects due to the location of rainfall gauges. Unfortunately, a large part of the natural variation in rainfall measurements is caused by a mix of spatial and temporal influences. Meteorological conditions are not spatially homogenous and orographic effects can depend on prevailing conditions. Importantly, exposure of rainfall gauges to an effect is generally dependent on meteorological conditions, primarily wind direction and speed. A ground-based rainfall enhancement trial was conducted using a randomized crossover design in South Australia in 2009. The analysis presented in this paper explores the limitations imposed by ignoring spatiotemporal variation in rainfall data and takes advantage of modern statistical methods to construct an appropriately specified model for these data. The level at which analysis is performed is addressed, particularly whether it is appropriate in this situation to use gauge-level data, as opposed to aggregated data such as average daily rainfall, in statistical inference. Our analysis, which accounts for the spatial and temporal correlation structure of rainfall data, suggest that there is a substantial increase in the likelihood of detecting a modification signal when the analysis is carried out at the gauge level.

1. INTRODUCTION

A prerequisite of rainfall enhancement detection is an estimate of the amount of rain which would have fallen had the treatment not been applied. In weather modification experiments, an estimate of this value is usually derived from mean rainfall measurements over the course of a trial within an untreated, control area. This estimate is compared with mean rainfall over a target area in an attempt to identify the degree of rainfall enhancement. A premise of this type of analysis is that a randomised treatment protocol, along with a sufficiently large sample size, will balance out the effects of variation between the target and control. However, very high levels of variability in rainfall make this assumption problematic. This variability is both temporal, across days, months and even years, as well as spatial, with some areas receiving a lot of rain while neighbouring areas concurrently receive little. Temporal variation in rainfall reflects the variation in meteorological conditions over time, while orographic factors such as elevation and proximity to a coastline will interact with meteorological conditions to produce spatially varied precipitation across adjacent areas.

This high level of variability produces a large spread of rainfall measurements, or large variance in the data. As a consequence, the calculated estimates of mean target and control rainfall are associated with large errors which is it necessary to reduce in order to obtain a significant result. This either requires a huge rainfall enhancement effect so that the consequent shift in mean rainfall is the main contributor to the observed variability in rainfall measurements, or a large enough sample size so that the control and target mean values are estimated precisely irrespective of the high variation in the actual rainfall measurements. The reality, however is that any rainfall enhancement signal is likely to be modest and the sample size limited, with the high variability of rainfall then creating a lot of background ‘noise’ against which it is often not possible to detect a rainfall enhancement signal. Various experimental design techniques can be employed to counter this variability, and thereby decrease the error in estimated means. By choosing a control area that is as similar as possible to the target area, one can minimise variation due to differences in the spatial distribution of rainfall measurements. This is the basic principle underpinning a typical crossover design, where two areas with similar rainfall patterns are selected and treatment is applied to either one or the other area using a randomised schedule. The untreated area acts as a control for the area being treated, and allows for a traditional double-ratio analysis. This approach goes some way to balancing random spatial variation. Increasing the sample size by running trials for

long periods, often many years, has been necessary to reduce the margin of error towards levels where analysis can potentially yield significant results. One can identify certain combinations of temporal conditions as providing a different level of rainfall enhancement from other combinations (known as ‘blocking’). Then by restricting randomised treatment allocation to blocks consisting of ‘similar’ temporal conditions, it may be possible to reduce the variance of outcome measures, thus reducing variability caused by day-to-day changes in these conditions. However, despite extensive efforts over the past fifty years to reduce the effect of natural variability in rainfall by employing these design techniques in randomised trials of rainfall enhancement, it seems clear this approach has not produced a reliable methodology for detecting rain enhancement (Cotton and Pielke 2007). Most cloud seeding experiments still do not provide evidence sufficient to reject the null hypothesis of no enhancement, and usually produce data ‘not sufficient to reach statistical conclusions’ (NRCNAS 2003).

The problem with this ‘pure’ randomisation-based analysis of rainfall is that it treats all rainfall measurements as being randomly drawn from the same distribution of potential values (in the sense that treatment application for these measurements was at random). However, we know that this is not true. We know that variation in rainfall from gauge to gauge within a day is driven by variation in orographic and meteorological conditions, and that day-to-day variation for a particular gauge is driven by variation in meteorological conditions. This constantly changing background of spatial and orographic variation typically has a much greater influence on rainfall than any potential enhancement effect. We can measure many of the important factors that potentially contribute to this background variation. However, this knowledge is ignored when we simply average target and control rainfall in the hope these external effects then cancel out. Insofar as they do not, the resulting variability in our estimation of an enhancement effect can (and typically does) far outweigh the actual level of this enhancement. This leads to significant limitations when a randomisation-based approach to evaluation of an enhancement effect is applied to this problem. More statistically efficient means of analysis are required if we hope to gain significant results in realistic time frames.

Statistical techniques have seen tremendous improvements in recent decades. However, within the field of weather modification, methodology is largely based on treatment randomisation and mean rainfall comparisons, and the use of alternative methods of statistical analysis have not, as yet, been widely adopted. Given the complexity of the problem at hand and the importance of any consequent improvements in water management, these significant

developments in statistical sciences ought to be explored by weather modification scientists. A review of statistical analysis in weather modification by the US National Research Council of the National Academies (NRCNAS 2003) makes the following assessment: "To fully consider and evaluate the myriad of variables in weather modification experiments, multivariate statistical process models that exhibit spatial and temporal dependence are much better suited [to this analysis than the single test]". Statistical methodologies for the analysis of spatiotemporal data were not available in the early days of weather modification. It was not until the 1990s that they could be realistically implemented in the analysis of field trials (see Cressie, 1991). Consequently randomisation has continued to serve to mitigate both spatial and temporal correlation effects in many evaluation studies. However, just as blocking designs can improve efficiency over randomisation, one can get more efficient estimates by modeling the spatial (and spatiotemporal) effects (NRCNAS 2003). The basis for these spatiotemporal methods is that evaluation is carried out in the context of statistical models that explicitly account for many of the numerous factors, which produce variation in rainfall. This is achieved by the inclusion of covariates, which are variables that influence rainfall (either directly or as proxies for unmeasured influences) but which themselves are not influenced by rainfall. Covariates are intended to control for variation in natural rainfall, making it easier to detect an enhancement signal if it exists. In contrast to a pure randomisation analysis, this type of analysis estimates the conditional contribution to rainfall by meteorological, orographic and enhancement factors. While potentially a much more powerful statistical technique than pure randomization, it is sensitive to choice of statistical model. A poorly specified model can lead to false inference.

Meteorological covariates that have some correlation with rainfall over both time and/or space include, but are not limited to, wind direction and speed as well as barometric pressure, air temperature and moisture content in the air. Important orographic covariates include the spatial location of a rain gauge relative to the source of the treatment and the elevation of the gauge. Gauge-specific orographic effects that are uncorrelated between gauges can be removed via the introduction of a random effect for each gauge in the model. These models can be estimated using maximum likelihood methods and the available daily rainfall data from each gauge (gauge-level data) within the target and control areas, rather than just the average rainfall measurements over all target and control gauges for a particular day. Overall, our experience is that gauge-level models that include the afore-mentioned meteorological and orographic covariates can account for around half the observed variation in gauge-level

daily rainfall (Beare et al. 2010). Consequently, if data for a large number of gauges are available (as is usually the case), then this approach has the potential to increase the power to detect an enhancement effect over a shorter time frame. It should be noted, however, that this extra power is dependent on our ability to adequately control for the between-gauge variability in rainfall on a day using appropriate spatial covariates in the analysis model. This information, which is averaged over when performing classical randomised, means-based, analyses, is effectively utilised in a model-based approach to derive a more accurate prediction of the natural rain in the target area, i.e. the rain that would have fallen in the absence of treatment. As a consequence, this yields a method to determine if a signal exists and a means to estimate the level of any effect. This is done by the taking the difference between the actual rainfall in the target gauges and the modeled estimate of the corresponding natural rain.

Accurate assessment of the precision of the estimated enhancement effect must also be considered. Given the complex nature of the statistical models for rainfall that underpins this estimate, standard approximation-based approaches to determining this precision, typically based on the use of t-statistics, become problematic as they often involve assumptions about rainfall measurements from individual gauges being uncorrelated. In this context adoption of a bootstrap approach to determining the distribution of the enhancement estimate, allowing direct calculation of confidence intervals for its expected value, is recommended. Bootstrapping involves creating an appropriate sampling design and repeated re-sampling of the trial data in order to determine the underlying correlation structure in the model errors that would overstate the precision of the estimate if these errors were assumed to be independent. The choice of sampling design is important and should reflect the likely sources of correlation such as spatial or temporal proximity. This is the approach taken for the model-based analysis described in this paper. In particular, we assess precision of the estimated enhancement via parametric bootstrap simulation based on the same models that are employed to obtain this estimate. We have applied these techniques to the analysis of a trial of a ground-based ionization rainfall enhancement technology known as Atlant, as an extension of the work in the 2008 trial of the technology, previously reported in the Journal (Beare et al. 2010). In addition we compare the modeling results and approach to the traditional double-ratio analysis.

Atlant is a ground-based ion generator that is deployed to increase rainfall downwind of the system. The physical processes by which it operates have not been established. However, it

has been hypothesized that the operation of the Atlant promotes the introduction of charged aerosols into the cloud layer, thereby increasing the coalescence of water droplets and hence promoting the formation of raindrops. The enhancing effect of electric forces on the collision/coalescence and formation of raindrops has been investigated experimentally since the 1950s (see Chapter 10 of McGorman and Rust, 1998 and Chapter 18 of Pruppacher and Klett, 1997 for an overview). However no substantive progress has been made in either identifying the underlying physical process or in the statistical measurement of the effect (if any) of such ionisation on rainfall.

2. DESIGN OF THE 2009 MOUNT LOFTY RANGES TRIAL

2.1 Experimental design

A primary aim of the 2009 Mount Lofty Ranges trial was to test the hypothesis that operation of the Atlant system in the assessment region lead to increased rainfall in the expected area of influence of Atlant. That is, there is a null hypothesis that corresponds to no effect on rainfall, with this being rejected in favour of evidence of rainfall enhancement if parameters (defining conditions) associated with operation of Atlant are significant in the statistical model for rainfall defined under this hypothesis, and this model also implies a statistically significant relative increase in observed rainfall that can be attributed to operation of Atlant. In taking this approach, it is recognised that the significance level used to test the hypothesis depends on the risk associated with falsely rejecting the null or alternative hypothesis. In order to test this hypothesis, it was decided to use a randomised crossover design with two operating sites. This allows the trial to be evaluated using the double-ratio analysis method as well as statistical modelling. Two distinct areas were chosen with the aim that they would be sufficiently near each other to expect a good correlation between natural rainfall in each area, but sufficiently far apart to reduce the chance of treatment in a target area influencing rainfall in a control area.


2.2 The Atlant locations and trial area

The Atlant site used in the 2008 Mount Lofty Ranges Trial at Willunga (C2) was again selected for the 2009 trial. In addition to this site, a second Atlant was introduced at Tea Tree Gully (C3) around 58 km north east of C2 (Figure 1). The Mount Lofty Ranges are orientated northeast to southwest, and both sites are situated along the first significant ridgeline closest to the coast, and are exposed to the prevailing weather—from the south west to north west. The elevations of C2 and C3 are 348 m and 373 m above sea level respectively, and have significant upslope valleys located to the west and northwest. At both sites, the landform elevation rises from the

coast travelling from west to east quite steeply up to 21% (21m vertical rise for every 100m horizontal run). An extensive rain gauge and weather station network is provided by the Bureau of Meteorology (BOM) within this trial area and includes 282 rain gauges. Of these, 262 gauges records of rainfall data over the period of the trial. The C2 and C3 sites are sufficiently near each other for there to be a reasonably strong correlation between the natural rainfall in each area (see Figure 1).

South Australia is classified as having a Mediterranean climate and is influenced by offshore trade winds in the summer and on-shore westerlies in the winter. As a consequence, the trial location experiences a dry and warmer period from November to April with prevailing winds from the southeast to east and a moderately wet and colder period from May to October with prevailing winds from the northwest to southwest (BOM, 2008). The climate of the Mount Lofty Ranges is significantly affected by an elevation ranging from 350 m to 700 m and winds sweeping across the Gulf of St Vincent. Typically, a moist marine onshore airflow from the west rises as it approaches these sites—i.e. there is orographic lifting. The resultant turbulence and vertical movement of air would be expected to result in upward dispersal of the any charged aerosols generated by Atlant, however the exact rate and dispersion of such particles has not been measured. Locating an external control area that matches the trial area is difficult. The meteorological and topographic characteristics of neighbouring areas were quite different from the trial area. The land area to the north and east of the trial area is relatively flat and dry when compared to the trial area, and the influence of offshore fronts on precipitation in these areas is not nearly as strong.

Figure 1 The location of the Atlant sites (Δ) at C2 (Willunga, $35^{\circ}18' 41.34'S$, $138^{\circ} 31' 22.02'E$) and C3 (Tea Tree Gully, $34^{\circ}49' 28.10'S$, $138^{\circ} 44' 48.70'E$). The rain gauges used in the trial are indicated by green dots. The circles centred on the Atlant sites have a radius of approx 90 km. Downwind sectors (yellow) are shown for a westerly wind. The degree of overlap is dependent on the direction of the wind.



2.3 Target area definition

The determination of whether a gauge is exposed to the ion plume generated by the Atlant technology cannot be ascertained with any degree of precision because our current understanding of the physics underlying the spread of the ion plume is incomplete. However, the plume will be directed by surface and upper level winds. The direction of these winds will change through the course of a day, but the typical range of wind directions over a 24-hour period is not great, with the analysis carried out for the 2008 trial at C2 indicating an average range of between 60° and 80° (Beare et al. 2010). It is therefore presumed that the path of a plume could spread across a similar sector originating at the Atlant device. While the selection of the angle of this sector is necessarily somewhat subjective, based on the wind direction variation in the 2008 trial, it seemed reasonable but conservative to deem this angle to be 60° . This angle was set before the trial commenced to reduce potential bias in the analysis. The orientation of this 60° sector is dynamically defined on a daily basis, being centred on the radial vector describing the downwind direction from the Atlant. Rather than including all gauges across a target area, this dynamic daily partitioning of the target area in relation to steering wind

direction serves to focus the signal generated by the Atlant system, if it is effective. The target areas for a westerly wind are shown in Figure 1. The determination of primary steering wind flows was obtained using radiosonde data (vertical wind profile), produced by the BOM. Radiosonde data was recorded at Adelaide Airport at six-hourly intervals to provide four readings per day. To obtain a steering wind direction, the vectors of wind speed and direction were averaged between the surface and upper winds (925 hPa and 850 hPa respectively) at each of the four time points throughout the day, and the vector average of these four averages was taken to obtain the daily wind direction. Further, a gauge was included in the analysis if it was within one degree of 'Euclidean distance' from one of the two Atlant sites. 'Euclidean distance' is the square root of the sum of the squared difference between the latitude of the gauge and the latitude of the site and the squared difference between the longitude of the gauge and the longitude of the site. This distance approximates 90 km.

2.4 Operating schedule

The trial ran for 128 days subject to the operating protocol described below, commencing at 9am 1 August 2009 and finishing at 9am 7 December 2009. All times were measured in local time. During the trial the Atlants were switched on and off at 9am in accordance with the specified switching regime. This was to coincide with the BOM reporting time for the rain gauges, and to reduce the chance that overlap of rainfall measurements diluted the results. An additional advantage is one of operational convenience, in that 9am is approximately the start of a working day. A 30-minute 'temporal buffer' was also added to the switch time, in recognition that there may be a delay, albeit of unknown length, between when the device is switched off or on and any effect on rainfall downwind of the device. Thus, with a nominal switch time at 9am, the operating Atlant was turned off at 8.30am and the ongoing Atlant was then turned on at 9am.

The two sites were operated according to a randomised, asynchronous, alternating daily schedule. C2 was operated on a randomised on-off sequence. C3 was operated on a randomised on/on-off/off sequence. Each of the four C2-C3 operating combinations (on-on, on-off, off-on, off-off) was scheduled for an equal number of days. Note that use of the above design was motivated by the need to compromise between maximising the cross-over frequency, which would have had C2 and C3 both operating on an asynchronous on-off basis, and the desire to allow testing for carryover Atlant operating effects from previous days, which required at least one site to operate continuously for two or more days. These carry-over effects had been noted in analysis of previous Atlant trials and while there was no physical reason to hypothesise that

such effects should occur, it was still important to test whether it represented something more than an artifact of the way the Atlant had been operated in those trials. In some instances an Atlant system was scheduled to operate but did not operate due to technical faults or in accordance with risk management plan. The Atlant operating data used in the preliminary analysis of the 2009 Mount Lofty Ranges trial reflected the actual duration of Atlant operation on a day. However, because there are a number of cases where either C2 or C3 ran or did not run for a very short time during the day, we adopted the rule that a site was treated as operational on a day if it ran for 12 or more hours that day. Otherwise it was treated as non-operational.

3. DYNAMIC DOUBLE RATIO ANALYSIS

3.1 Double-ratio analysis

Statistical methods based on the double-ratio statistic have been widely used in the field of weather modification, and are typically applied to rainfall data to determine the effect of cloud seeding (Gabriel, 1999). Comparison of rainfall over seeded versus non-seeded periods allows inferences on the effects of cloud seeding to the extent that all differences between the two periods can then be ascribed either to the effect of seeding or to random year-to-year variation. The double ratio statistic has an expected value of one if seeding has no effect and there is evidence of a positive effect of seeding upon rainfall if its value is significantly greater than one.

In the crossover design used for the trial, paired target areas C2 and C3 are 'seeded' (Atlant turned on) on an alternating basis, with the 'unseeded' area (Atlant turned off) serving as the control for the seeded target area. That is, the downwind area of C2 acts as a control for the downwind area of C3 when C2 is off and C3 is on, and *vice versa*. The effect of seeding (Atlant operation) can then be assessed using the value of the root double ratio (RDR), which is the geometric mean of the ratios of the area-specific seeded to unseeded precipitation. A requirement of double-ratio analysis is that the area downwind of C2 does not overlap with the downwind area of C3, and *vice versa*.. However in the 2009 trial, overlaps between target and control areas occurred. In addition, these areas are dynamically defined, that is defined differently from time to time, depending on the prevailing meteorological conditions. Hence a variant of double-ratio analysis was carried out for the trial period. This is called a dynamic

double-ratio (DDR) statistic, because it reflects that the target and control areas were redefined each day depending on wind direction and did not overlap.

The formula for the DDR used for the 2009 trial is:

$$DDR = \sqrt{\frac{\text{Downwind C2 only}_{C2on/C3off}}{\text{Downwind C3 only}_{C2on/C3off}} \times \frac{\text{Downwind C3 only}_{C2off/C3on}}{\text{Downwind C2 only}_{C2off/C3on}}}$$

where $\text{Downwind C2 only}_{C2on/C3off}$ denotes the ‘average rainfall’ recorded by gauges that were downwind of C2 but not of C3 on days when C2 was operational but C3 was not. Similar interpretations hold for the other components of DDR.

There are a number of ways that the DDR can be calculated, depending on how one interprets the concept of ‘average rainfall’. These include:

- The average of all gauges over all days (gauge by day average) of total observed rainfall in all gauges in each of the four downwind areas defined in the DDR over the period of the trial;
- The simple average of daily average rainfall in each downwind area; and
- The area weighted average of daily average rainfall in each downwind area, where the area used for each gauge by day rainfall reading is the area of the Voronoi polygon centred at the gauge that provided the reading. This is the polygon defined by locations surrounding the gauge that are closer to it than they are to any other gauge in the trial area.

The first DDR definition is, from a statistical perspective, the most efficient but it does give greater weight to days on which there were more observations at individual gauges. Further, in determining the accuracy of the estimate it would be necessary to take into account the spatial correlation between the gauge observations. The second definition gives equal weight to all days and can be seen as a comparison of estimates of average daily rainfall in the downwind areas. However, it does not take into account the fact that the spatial distribution of gauges is far from uniform. Weighting by the area for which a particular gauge is the closest observation, i.e. its Voronoi area, is an attempt to allow for this.

3.2 Results of the DDR analysis

The DDR statistics for the trial under each method are presented in Table 1. There is an implicit assumption in the preceding Section that there are two downwind areas from which rainfall data

are obtained for each day when C2 is on and C3 is off and vice versa. Out of the 128 days of the trial, there were 63 ‘on-off’ days when this condition held. However, only 46 of these provided rainfall data for both downwind areas used in the DDR statistic. Consequently, in Table 1 below we present average rainfall values and the resulting DDR value under each of the three methods described above when data from all 63 ‘on-off’ days are used and also when only data from the 46 ‘balanced on-off’ days are used in the calculation of the DDR statistic. We note that the DDR statistics defined by averaging across all gauge-day records, or by weighted averaging of these records based on the area of the Voronoi polygon surrounding a gauge. DDR’s based on the Voronoi weighted averages have values greater than one, while those based on simple averaging of daily average rainfall have values less than one. These differences highlight the sensitivity of the DDR statistic to instability in the values of daily average rainfall, and indicate the need for a more sophisticated modeling approach that does not rely (as does the DDR analysis) on the success of randomisation across days in balancing meteorological and orographic influences on rainfall.

Table 1 DDR statistics for 2009 trial.


Component	C2On/C3Off Downwind C3 (C2 Control)	C2On/C3Off Downwind C2 (C2 Target)	C3On/C2Off Downwind C2 (C3 Control)	C3On/C2Off Downwind C3 (C3 Target)	DDR value
All 63 ‘on-off’ days					
# gauge-day records	802	774	1187	691	
Overall Average	2.9868	1.8160	1.0268	2.6107	1.2434
Average of Daily Averages	3.2954	2.1102	1.8715	2.3090	0.8888
Voronoi-Weighted Averages	3.0015	2.4347	1.3522	2.6749	1.2668
46 ‘balanced on-off’ days					
# gauge-day records	627	238	449	643	
Overall Average	3.8188	3.2361	1.1506	2.8056	1.4375
Average of Daily Averages	3.8936	2.3403	2.1349	2.4052	0.8229
Voronoi-Weighted Averages	3.4462	3.5701	1.5468	2.7923	1.3676

The fundamental question is how significant might any of these estimated DDRs be. One obvious concern relates to the issue of spatial correlation in gauge-level observations. If it is reasonable to expect that gauges roughly oriented in a downwind direction are more likely to be correlated then this could, in turn, lead to an understatement of the variation in the daily averages and the inter-annual variation in total observed rainfall. This could easily occur if the spatial variation in rainfall on a given day was more homogenous downwind than crosswind. To address this question more broadly we examined a set of hypothetical DDRs using historical data.

3.3 Historical DDR analysis

In order to examine the underlying variability in the DDR over time, we look at the inter-annual variation based on overall downwind area averages, as this should be the most efficient estimate. The DDR values were calculated from historical data using the same operating schedule and site locations as for the 2009 trial. The analysis is based on daily gauge-level rainfall over the trial area for August–December 1998–2007. The year 2008 was excluded, as there was an Atlant trial in the trial area over the same period. The results are summarized in Figure 2 and also plotted are 2009 statistics from Table 1 for comparison. From the figure it can be seen that three out of the 10 historical observations exceeded the trial result, and it also highlights the high level of year-to-year variability in the DDR statistic. While the DDR is a simple and easy to calculate statistic for an experiment with a randomised crossover design, it is clear that it does not provide the necessary level of control to detect an enhancement signal that is small relative to the background variation in natural rainfall. Given the lack of stationarity of rainfall patterns due to medium to long-term variations in climate it seems unlikely this problem could be overcome though extended trials. While there are undoubtedly a number of improvements that would be made to the design of the trial, it is inherently difficult to guard against a large number of meteorological and orographic factors that affect rainfall and for which there is a considerable degree of uncertainty. Hence in the following analysis we have elected to pursue a conditional approach that estimates Atlant effects using regression models based on meteorological and orographic covariates.

Figure 2 Historical DDR statistics based on gauge-day average rainfall, 1998–2009.



4. GAUGE-LEVEL ANALYSIS

4.1 Gauge-level rainfall correlation

The statistical analyses of previous Atlant trials were conducted using gauge-level data as the units of analysis. In contrast to using daily averages of gauge-level rainfall across an area, the use of gauge-level data allows for exploration of interactions between meteorological conditions and the orographic effects associated with individual gauge locations. Taking advantage of the information available at each gauge by day level reduces the background noise against which a rain enhancement signal can be sought. This approach, however, presents the analytical challenge of properly accounting for the correlation in rainfall values recorded by different gauges on the same day. Standard errors for the coefficients of an ordinary least squares regression model are based on the assumption that the observations are mutually uncorrelated. Where such correlation exists, the effect is to increase the standard errors associated with the estimates of the model parameters. This can seriously impact on the validity of statistical inferences for these parameters, since test statistics can be biased upwards, leading to claims of significance where in fact no such significance actually exists.

The degree to which rainfall at a gauge is correlated with that of its neighbour is not straightforward to measure and is not constant over time. In a spatial sense, whether a gauge is measuring the same thing as its neighbour or something different will depend on both the distance to its neighbour and the geographical extent of the weather event being measured. In a temporal sense, weather systems exhibit strong cyclical behavior that will tend to link current and past observations of meteorological conditions. The challenge is to achieve an effective measure of this correlation. At one extreme, if all gauges were perfectly correlated we would have, in effect, only one observation. In fact, the day-level analyses presented earlier represents this extreme in that they effectively assume that the correlation between gauges downwind of either Atlant site is perfect on any given day, and so we have exactly one observation per day for the gauges downwind of C2 and another for gauges downwind of C3. Taking into account the possible pair-wise correlation between these two averages then gives us a reasonable idea of the upper bound of the relative standard deviations of the model coefficients. In contrast, a gauge-level analysis performed under the assumption that rainfall measurements are mutually uncorrelated from day to day and from gauge to gauge provides a lower bound for these values. This situation can be improved upon by the specification of an appropriate model for this correlation and by the use of bootstrapping techniques based on this model for statistical

inference. We expect that observations that are nearby or closely related exhibit a greater degree of correlation. The question then becomes one of what is meant by nearby or closely related. Gauge-level rainfall may be correlated in a number of dimensions, including: (1) Distance—correlation increases as distance between gauges declines; (2) Elevation—correlation increases as the absolute or vertical difference in elevation between gauges declines; and (3) Time—correlation increases as observations become closer in time.

Correlation structure may also emerge through more complex processes. For example, the spatial correlation between gauges may be greater in an upwind-downwind direction as opposed to a crosswind direction. As wind direction shifts over time the spatial correlation structure will shift. When rainfall is modelled as a function of meteorological and orographic covariates this removes part but not necessarily all of the correlation structures between gauge-level rainfall observations. Our strategy, in terms of being able to draw statistical inferences with regard to the efficacy of Atlant, is therefore two-fold. First, construct a set of covariates that remove as much as possible of the spatio-temporal correlation in gauge-level rainfall; and second, use random effects to correct for any remaining spatio-temporal correlation, thus ensuring that reported standard errors are no longer biased downwards because of a lack of mutual independence. However, there is still a substantial degree of subjectivity involved in implementing this strategy. In a single dimension with evenly spaced observations it is straightforward to enumerate the structure of any existing correlations. In multiple dimensions with unevenly spaced observations, such an implementation strategy has to be based on a subjective assessment of the potential but unmeasured sources of correlation between rainfall measurements.

4.2 Structure of the gauge-level model

The gauge-level modeling is done in two stages. First, the probability of observing positive rainfall, a rainfall event, at a gauge is modeled using a logistic regression. In the second stage the expected level of rainfall, given that a rainfall event has occurred, is modeled. Here we model the natural logarithm of rainfall for reasons discussed in the previous section. One reason for using a two-stage approach is that the logarithm of zero is not defined. A second reason is to correct for the bias that occurs in a single stage regression model when a substantial number of observations are clustered at a single point, in this case zero rainfall. By combining the estimates obtained from these two stages we can make an unbiased estimate of the expected level of rainfall that would be recorded at a given rainfall gauge on a given day. This estimate is conditional on the observed values of a number of covariates that influence rainfall directly (or

are a proxy for unobserved factors that influence rainfall) and are themselves, not influenced by the level of precipitation falling at a given time and location. The covariates also include variables that represent the operating status of the two Atlant systems or serves as proxies for the extent to which a gauge is exposed to an actively operating site. See Table 2 for a summary of the categories of covariates used.

Table 2 Summary of categories of covariates used in spatiotemporal models.

Covariate Type	Description
Meteorological	Wind direction and speed at each of three altitudes (as opposed to an average across two altitudes), measured at 700, 850 and 925 hPa; Lagged (previous day) meteorological observations, denoted with an L1 suffix.
Gauge-specific orographic	Gauge elevation; Gauge distance from the operating sites.
Atlant operating sequence	Dummy variables that identify a site's operating status on the day as well as on the previous day. These account for differences in the one-day lag structure of the operating schedule for C2 and C3. This is necessary because C2 was operated on a randomly assigned 2-day cycle (on-off), while C3 was operated on a randomly assigned 4-day cycle (on/on-off/off).
Gauge-specific spatio-temporal	Gauge location relative to the steering wind direction at each site on a day as well as on the previous day; Gauge distance from each site relative to the operating status of the site on that day as well as on the previous day.

4.2.1 Random effects


Random effects are typically used to allow for correlations in the data that are not captured by the model covariates. Typically we think of these correlations as being either spatial or temporal. For example, a purely spatial random effect that allows for potentially unmeasured orographic variables that varies from one gauge location to the next. However, given the location of the gauges in the target area changes from day to day according to the steering wind direction, it would be reasonable to postulate that there would be a mixed spatio-temporal correlation structure between gauge level observations. This could be quite strong if, for example, rainfall patterns tended to track more closely in a downwind as opposed to a cross wind basis. We can model such a spatio-temporal random effect by grouping gauge-level observations on a day according to their downwind and crosswind orientations relative to the two Atlant sites. Note that inclusion of this spatio-temporal effect reflects the hypothesis that the correlation between gauge-level rainfall measurements on a day will be stronger in a downwind direction on that day as opposed to a crosswind direction. The definition of the spatio-temporal groups was subjective, but aimed at grouping downwind gauges that were likely to have had similar exposure to prevailing meteorological conditions, including exposure to Atlant. They were constructed in terms of the radial angle made by a gauge with C2 relative

to the average steering wind direction for the day ($C2\theta$) or the same angle but now relative to $C3$ ($C3\theta$). Six radial classes were constructed as listed in Table 3. For each day, these classes effectively form a radial tiling of the downwind region at each site (see Figure 3).

Table 3 Definition of radial classes for spatio-temporal random effects.

a	$C2\theta$ greater than 30° and $C3\theta$ greater than 60°
b	$C2\theta$ less than or equal to 30° and $C3\theta$ greater than 60°
c	$C2\theta$ less than $C3\theta$ and both $C2\theta$, $C3\theta$ less than or equal to 60°
d	$C3\theta$ less than or equal to $C2\theta$ and both $C2\theta$, $C3\theta$ less than or equal to 60°
e	$C3\theta$ less than or equal to 30° and $C2\theta$ greater than 60°
f	$C3\theta$ greater than 30° and $C2\theta$ greater than 60°


Figure 3 Areas covered by radial classes in a north-west wind.



Theoretically, this tiling leads to $128 \times 6 = 744$ spatio-temporal groups. However, many of these are empty, leading to 664 such groups containing rainfall data over the period of the trial. They account for a sizable component of the variability in rainfall. The distribution of group sizes, defined by the number of gauge-day records with non-missing rain, is shown in Figure 4. The mean group size is 13.2, with a median of 10. Note however, that the average group size drops to 4.8 and the median group size drops to 1 if group size is defined in terms of records with positive rainfall. There are 397 spatio-temporal groups with at least one positive rainfall reading. The large number of spatio-temporal groups and the way in which they were constructed again leads to the possibility that the random effects will attenuate part of the Atlant signal. The radial angle and distance to the mid-point of each cluster may be a proxy for

exposure to the Atlant ion plume on a given day. This would lead to an under-attribution if there is a positive Atlant contribution to rainfall. This potential loss of power for detecting an Atlant signal was viewed as acceptable given the increased precision in rainfall modelling due to inclusion of the spatio-temporal random effects.

Figure 4 The size (number of gauge-days with non-missing rain) distribution of the 664 groups used to define the spatio-temporal random effects.



4.3 Probability of a rainfall event

The probability of a rainfall event at a given gauge location on a particular day was modelled using a logistic specification. Note that this model is fitted under the assumption that gauge-level rain events are independently distributed day to day and gauge to gauge. A logistic model with spatio-temporal random effects was also fitted using a Laplace approximation to the likelihood (via the *glmer* function in the R statistical package). However, the fitting procedure failed to converge. It was therefore decided to proceed with a standard logistic model fit based on an independence assumption. As a consequence the significance of the Atlant and other model covariates in the model are potentially overstated. The logistic model fit for a rainfall event defined in terms of a threshold of 0.1 mm rainfall in a gauge is presented in Table 4. All covariates in the model are real-valued (including the dummy variables), with interactions defined by straightforward multiplication of the corresponding main effects. Note that a positive sign for a model coefficient indicates that an increase in the value of the corresponding covariate increases the probability of a rain event. Of the eight coefficients in the model associated with operation of the Atlant (the last eight in Table 4), only three are significant. Two of these are marginal (p-values greater than 0.04), with the third (p-value = 0.02) associated with the interaction between distance from C2 and its operational status. Given the potential overstatement of significance due to the assumption of gauge-level independence of rain events, it therefore seems reasonable to view the impact of Atlant on the occurrence of rain events as

negligible. However, the model presented in Table 4 is still useful because as it allows us to use bootstrap simulation to generate alternative rainfall data that include the observed spike at zero in the gauge-level rainfall record. This spike arises because, on a day to day basis, most gauges do not record positive rainfall.

Table 4 Logistic regression fit for the probability of downwind gauges recording a rainfall event of 0.1mm or more. Based on 8661 downwind gauge-days of reported rainfall.

Parameter	Estimate	Standard Error	P-Value
Intercept	0.6352	9.9020	0.9489
August/September	-0.1700	0.1208	0.1592
WRE	0.4684	0.3288	0.1543
Upwind Rainfall	5.7195	0.1974	0.0000
Wind Speed 700	-0.0152	0.0035	0.0000
Wind Speed 700 - L1	0.0177	0.0034	0.0000
Wind Speed 850	0.0362	0.0073	0.0000
Wind Speed 850 - L1	-0.0114	0.0062	0.0660
Wind Speed 925	-0.0089	0.0068	0.1931
Wind Speed 925 - L1	0.0003	0.0053	0.9509
Wind Direction 700	-0.7120	0.4543	0.1171
Wind Direction 700 - L1	-0.4384	0.3865	0.2567
Wind Direction 850	1.2205	0.6971	0.0800
Wind Direction 850 - L1	-0.2299	0.6104	0.7064
Wind Direction 925	-0.3616	0.4260	0.3960
Wind Direction 925 - L1	-1.2559	0.3701	0.0007
Air Temperature	-0.0929	0.0224	0.0000
Dew Point Depression	-0.0175	0.0190	0.3556
Sea-level Pressure	-0.0007	0.0097	0.9432
Elevation	0.2534	0.0252	0.0000
Distance C2	-1.5233	0.7100	0.0319
C2 θ	0.0007	0.0018	0.7188
Distance C2 * C2 θ	0.0065	0.0034	0.0584
C2 θ - L1	-0.0035	0.0013	0.0048
Distance C2 * C2 θ - L1	-0.0006	0.0032	0.8575
Distance C3	-1.0923	0.7229	0.1308
C3 θ	-0.0005	0.0009	0.5663
Distance C3 * C3 θ	-0.0021	0.0033	0.5183
C3 θ - L1	-0.0016	0.0009	0.0755
Distance C3 * C3 θ - L1	0.0115	0.0032	0.0003
C2 On	-0.4163	0.2101	0.0475
Distance C2 * C2 On	0.6550	0.2842	0.0212
C2 On - L1	-0.4223	0.2086	0.0429
Distance C2 * C2 On - L1	-0.0476	0.2727	0.8613
C3 On	-0.0184	0.2028	0.9277
Distance C3 * C3 On	0.0908	0.3914	0.8166
C3 On - L1	-0.2142	0.1637	0.1905
Distance C3 * C3 On - L1	-0.4286	0.3012	0.1548

4.4 The *Log Rainfall* model

A linear model with spatio-temporal random effects for *Log Rainfall* was fitted using restricted maximum likelihood, with the resulting fit shown in Table 5. The fixed effects covariates in the model are the same as those set out in Table 2. The percentage column in the variance components of the random effects shows the proportion of the unexplained variation in *Log Rainfall* that is associated with each random effect in the model. Table 5 also shows the estimated Atlant attribution, both as a percentage of actual rain as well as estimated natural rainfall over the trial. The method for calculating the Atlant attribution is the same as that described in the analysis of the 2008 Atlant trial (Beare et al, 2010) and involves estimation of the level of natural rainfall at a gauge (i.e. the rainfall that would have been expected at that gauge if one or both Atlants had not been operating that day or the day before).

With random effects models, the issue arises of what to do with the predicted random effects (the group or level two residuals) when estimating the Atlant contribution. Three simple alternatives include: (1) Apportion these residuals to natural rainfall; (2) Apportion these residuals to observed (augmented) rainfall; and (3) Apportion these residuals in proportion to the ratio of natural to observed rainfall. There is no correct choice. Given these predictions sum to zero there is no inherent overall bias although there may be regional bias if the level two residuals are spatially correlated as the estimates are done at the gauge level. The first alternative above is likely to be the most conservative approach, and so was adopted.

The model parameter estimates are set out in Table 5. The overall fit of this spatio-temporal random effects model accounts 71 per cent of the variation in gauge-level rainfall, with the spatio-temporal effects accounting for almost 32 per cent of the observed variation in gauge-level rainfall. The estimated standard errors are inflated when compared with a model that only allows for purely spatial random gauge effects, reflecting the presence of significant within-day correlation for gauge-level rainfall measurements. Notwithstanding, five out of the last eight covariates in the model (i.e. those associated with operation of Atlant) are significant.

Table 5 Log Rainfall regression results with random spatio-temporal effects. Based on 3177 observations of positive gauge-level rainfall. Fitted values (including predicted spatio-temporal effects) account for 71 per cent of overall variation in Log Rainfall.

Parameter	Estimate	Standard Error	P-Value
Intercept	29.7520	9.9115	0.0027
August/September	-0.5116	0.1188	0.0000
WRE	1.0761	0.1551	0.0000
Upwind Rainfall	2.3431	0.1841	0.0000
Wind Speed 700	-0.0049	0.0036	0.1737

Wind Speed 700 – L1	0.0089	0.0033	0.0070
Wind Speed 850	0.0032	0.0069	0.6395
Wind Speed 850 – L1	-0.0018	0.0059	0.7578
Wind Speed 925	0.0070	0.0064	0.2763
Wind Speed 925 – L1	-0.0057	0.0051	0.2645
Wind Direction 700	-0.6828	0.4688	0.1461
Wind Direction 700 – L1	0.0447	0.4297	0.9171
Wind Direction 850	-0.4951	0.7340	0.5004
Wind Direction 850 – L1	0.2796	0.6176	0.6510
Wind Direction 925	0.4223	0.4283	0.3248
Wind Direction 925 – L1	-1.2170	0.3549	0.0007
Air Temperature	-0.0243	0.0243	0.3170
Dew Point Depression	-0.0128	0.0211	0.5448
Sea-level Pressure	-0.0295	0.0097	0.0026
Elevation	0.1410	0.0124	0.0000
Distance C2	0.1559	0.4626	0.7361
C2 θ	0.0032	0.0011	0.0033
Distance C2 * C2 θ	-0.0004	0.0021	0.8552
C2 θ – L1	-0.0020	0.0008	0.0093
Distance C2 * C2 θ – L1	-0.0026	0.0020	0.1888
Distance C3	-0.7602	0.4789	0.1125
C3 θ	-0.0013	0.0005	0.0085
Distance C3 * C3 θ	0.0042	0.0021	0.0434
C3 θ – L1	-0.0010	0.0005	0.0248
Distance C3 * C3 θ – L1	0.0088	0.0019	0.0000
C2 On	0.4523	0.1448	0.0019
Distance C2 * C2 On	-0.5654	0.1608	0.0004
C2 On – L1	0.1176	0.1384	0.3961
Distance C2 * C2 On – L1	-0.4978	0.1570	0.0015
C3 On	-0.0608	0.1593	0.7029
Distance C3 * C3 On	0.5149	0.2307	0.0257
C3 On – L1	0.2863	0.1141	0.0126
Distance C3 * C3 On – L1	-0.1463	0.1504	0.3308
Variance Component Estimates			
Random Effect	Component	Percentage	
Spatio-temporal	0.30587	31.870	
Residual	0.65386	68.130	
Atlant Attribution			
Actual Rainfall (mm)	13640		
Estimated Natural Rainfall (mm)	12465		
Attribution (% actual rainfall)	8.6		
Attribution (% natural rainfall)	9.4		

4.5 A semi-parametric bootstrap analysis of Atlant attribution

While it is difficult to interpret the coefficient estimates from the models defined in Table 4 and Table 5, they can be used to calculate the effect of changes in the operating status of an Atlant system on both the probability of observing a rainfall event and the expected level of rainfall given a rainfall event occurs. These quantities in turn can be used to derive gauge-by-day estimates of the change in rainfall that can be attributed to the operation of the Atlant system. Note that these estimated changes can be positive or negative, so it is their total over the period

of the trial that is of interest. The method used to calculate this attribution statistic is detailed in Section 5 of Beare et al (2010). It is sufficient to say here that it is a complex statistic, whose distribution depends in a nonlinear way on the regression coefficients and the variance components (the variances of the random effects) of the underlying linear mixed model.

Standard normal theory methods for constructing confidence intervals for complex statistics can be sensitive to the assumption of underlying normality. This is especially the case when the complex statistics are based on estimated random effects, as is the case for the attribution estimates presented in Table 5. Ideally, one would like to simulate or 'bootstrap' the distribution of these attribution estimates in as non-parametric way as possible in order to get around this problem. However, fully non-parametric bootstrap methods for data containing random effects are not well developed, and so we adopt a semi-parametric block bootstrap approach instead. A block bootstrap simulates the sampling distribution of a complex statistic based on correlated data by resampling blocks of data values rather than individual data values, with the blocks constructed so that, as far as possible, they contain individual values that are correlated with one another within a block, but are uncorrelated across blocks. In the context of the linear mixed model with spatio-temporal random effects underpinning Table 5, the blocks are the spatio-temporal groups underpinning these effects, i.e. the daily varying radial classes defined in Table 3, and the values associated with each block are its average residual and the within-block deviations from this average that together make up the unconditional residuals for that block under the model defined in Table 5.

In order to construct the semi-parametric block bootstrap distribution for a statistic that depends on the gauge-level models of Table 4 and Table 5. The logistic model set out in Table 4 was first used to simulate rain events for all 8661 gauge-days contributing to this table. Working only with blocks containing gauges designated as recording a rain event following this process, two independent random block samples were selected with replacement, with the first sample contributing the average residuals for these blocks, and the second contributing the within-block deviations for all gauges identified as recording rain in the same blocks. These values were then combined with the estimated fixed effects generated via the model in Table 5 to produce a set of simulated log rainfall data values for each such block. Finally, these log values were exponentiated in order to recover actual rainfall measurements.

In order to ensure that the rainfall distribution generated in this way is as realistic as possible, three further modifications were made to this distribution - all simulated rainfall values of less than 1 mm were rounded to the closest multiple of 0.2 mm; all simulated daily rainfall values

greater than 50 mm were randomly restricted to lie between 25 mm and 50 mm; and the total simulated rainfall for the trial period was randomly restricted to lie between 10,000 mm and 17,000 mm (the actual total for 2009 was 13,640 mm). Finally, these simulated rainfall values were used to refit the models in Table 4 and Table 5, and all statistics (including attribution estimates) were re-calculated. The procedure outlined in the preceding paragraph was independently carried out a large number of times in order to generate bootstrap distributions for the statistics of interest. For the results quoted below we used 10,000 bootstrap repetitions.

A known problem with use of a non-parametric block bootstrap where the blocks correspond to groups in a mixed model is that the resulting bootstrapped estimates of the variance components in this model are negatively correlated. This can lead to substantial undercoverage of their resulting bootstrap confidence intervals. In order to overcome this problem, a post-bootstrap correction was made to the bootstrap distributions obtained using the steps described above. First, the multivariate bootstrap distribution of the variance component estimates for the model defined by Table 5 was transformed ('tilted') in order to ensure that these estimates are uncorrelated. Next, all bootstrap distributions of model parameter estimates were 'tethered' to the original estimate values, using either a mean correction (for estimates, e.g. regression coefficients, defined on the entire real line) or a ratio correction (for estimates, e.g. variance components, that are strictly positive). Finally, all bootstrap distributions of complex statistics (including attribution estimates) dependent on these parameters were recomputed. Note that these distributions were also 'tethered' to their original sample values. The attribution estimates set out at the bottom of Table 5 are unweighted, reflecting only the estimated increase in level of rainfall as measured in the downwind gauges that contributed to the modelling process. Since these gauges are not distributed uniformly over the trial area, the attribution estimates in Table 5 should not be interpreted as estimating the increase in the volume of rain that fell in this area. Accurate conversion of rainfall gauge level measurements to volume of rainfall on the ground requires sophisticated spatial modelling and prediction that is beyond the scope of this report. However, a crude estimate of rainfall volume can be obtained by multiplying each gauge rainfall reading by the area of the Voronoi polygon surrounding the gauge and then summing over gauges. Note that a Voronoi polygon for a particular gauge identifies the region containing points that are all closer to that gauge than they are to any other gauge. Consequently, attributions were calculated on an unweighted and on a Voronoi area weighted basis, with the latter providing estimates that are more closely aligned with the volume of rainfall that fell in the trial area. Note, however, that Voronoi weighting tends to give large weights to gauges in

regions with sparse coverage. As consequence, the weighted attribution estimates tend to be relatively more variable.

In summary, the unweighted Atlant attribution estimate (as a percentage of estimated level of natural rainfall) was:

- An average enhancement effect of 9.4 per cent; with
- A semi-parametric block bootstrap standard error of 6.8 per cent ($p = 0.07$).

In contrast, the corresponding Voronoi weighted Atlant attribution estimate was:

- An average enhancement effect of 5.6 per cent; with
- A semi-parametric block bootstrap standard error of 6.4 per cent ($p = 0.19$).

A key output from the bootstrap analysis is the capacity to measure of the confidence in the Atlant attribution being positive. In this context, Table 6 presents bootstrap-based one-sided confidence intervals at different confidence levels for the unweighted and Voronoi weighted attributions (as a percentage of estimated natural rainfall) over the 2009 trial. The corresponding bootstrap distributions of these attribution statistics are shown in

Figure 5 and Figure 6. It can be seen that from an unweighted perspective we are 90 per cent confident that there was a positive Atlant effect over the trial period. This drops to 80 per cent when we move to a perspective based on Voronoi weighted or volumetric attribution. This decline is not unexpected given that large weights are given to gauges in sparsely covered areas, and reflects the fact that there is inherently more variability in volumetric rainfall measurements than in gauge measurements when gauges are widely scattered over a large area. As a consequence the precision of any volumetric estimate of rainfall based on Voronoi weighting of gauge data may not be very high.

Table 6 Lower bounds for parametric bootstrap estimates of one-sided confidence intervals for Atlant attribution based on level (unweighted) and volume (Voronoi area weighted). Note that % relates to corresponding estimates of natural rainfall.

Confidence level	99%	95%	90%	80%	70%	60%	50%
	Unweighted						
Atlant attribution (%)	-4.8	-1.1	1.0	3.6	5.6	7.4	9.1
	Voronoi area weighted						
Atlant attribution (%)	-8.0	-4.2	-2.3	0.1	2.1	3.7	5.3

Figure 5 Bootstrap distribution of Atlant attribution (as a % of estimated level of natural rainfall) under model with random spatio-temporal effects, unweighted. Blue curve shows a smooth non-parametric density fitted to the bootstrap values of this statistic, while red dashed line shows estimate obtained using actual *Log Rainfall* data.





Figure 6 Bootstrap distribution of Voronoi-weighted Atlant attribution (as a % of estimated volume of natural rainfall derived using the same Voronoi weights) under model with random spatio-temporal effects. Blue curve shows a smooth non-parametric density fitted to the bootstrap values of this statistic, while red dashed line shows estimate obtained using actual *Log Rainfall* data.



5. CONCLUSIONS FROM THE TRIAL ANALYSIS

The gauge-level analysis of the 2009 trial data focused on the models that allowed for spatio-temporal effects in the rainfall data. A particular type of spatio-temporal correlation structure was assumed that appears consistent with how steering winds are thought to affect the spatial distribution of rainfall. The analysis showed an overall enhancement effect of 9.4 per cent, based on a simple average gauge-level attribution, over the course of the trial. Based on an area-weighted gauge average the attribution was 5.6 per cent. This is relative to the 'natural' rain that would otherwise have been observed. The analysis also indicated that operation of Atlant had a negligible effect on the probability of rainfall being observed at a gauge.

Finally, the issue of a proper assessment of the variability of the estimated Atlant enhancement was dealt with via spatio-temporal bootstrap simulations. These simulations showed that the simple average attribution was significant at over the 90 per cent confidence level, and the area-weighted attribution was significant at over the 80 per cent confidence level. The loss of accuracy in using the area-weighted averages was not unexpected given the large weights that were given to a limited number of isolated gauges located at the outer edges of the trial area.

The gauge-level analysis strongly indicated that a spatio-temporal modelling approach can more effectively identify a signal in weather modification experiments when compared to traditional double-ratio tests. Further, semi-parametric bootstrap methods can be used to account for the correlation between gauges and provide a robust assessment of the precision with which rainfall enhancement effects can be measured.

6. REFERENCES

- Beare, S., R. Chambers, S. Peak, 2010: Statistical Modeling of Rainfall Enhancement. *Journal of Wea. Modif.* 42, 13-32.
- BOM (Bureau of Meteorology) (2008). Australian Bureau of Meteorology, Learn about meteorology, Climate Education, Climate of Australia, [Online Accessed on 9 November 2008] URL: <http://www.bom.gov.au/lam/climate/levelthree/ausclim/ausclimsa.htm>
- Cotton, W.R. and R.A. Pielke Sr., 2007: *Human Impacts on Weather and Climate* 2nd ed.
- Cressie, N. A. C. 1991, *Statistics for Spatial Data*, Wiley, New York.
- Gabriel, K. Ruben, 1999: Ratio Statistics for Randomized Experiments in Precipitation Stimulation. *J. Appl. Meteor.*, 38, 290–301.
- McGorman, D.R. and W.D. Rust: *The Electrical Nature of Storms*. Oxford University Press. 1998. 422pp
- National Research Council of the National Academies of Sciences (NRCNAS), 2003: *Critical Issues in Weather Modification Research*. National Academies Press pp41.
- Pruppacher, H.R. and J.D. Klett. 1997: *Microphysics of Clouds and Precipitation*. Kluwer Academic Publishers. 954pp