capstract:2611.00041v1 [astro-ph.IM] 20 Nov 2026

Cloud onset following scheduled completion of an observing sequence

Mira Hale,[1] Nadia Venn,[1] and Tomas Rell[2]
[1]*Department of Astronomy, Calwick University, Valen, Canada*
[2]*Estmere Observatory, Institute for Optical Astronomy, Valen, Canada*
(Dated: November 20, 2026)

Cloud forecasts for telescope scheduling normally use measurements of the sky without reference to the observing programme. We tested whether the scheduled completion of a sequence affects short-term cloud onset at a 1.2-m optical telescope. On 400 clear-start nights we independently assigned the operator to stop or continue after a common 15-min sequence and to announce either that the sequence was the last or that it was beginning. The programmed telescope commands during the measurement window were identical. A fixed-sector all-sky camera detected cloud onset on 91 of 200 stop nights and 36 of 200 continue nights, a difference of 27.5 percentage points (approximate 95% interval 18.8 to 36.2). The marginal difference associated with the sentence was 2.5 points (interval -6.6 to 11.6). The optical rule agreed with blinded optical and infrared review on 189 of 200 selected sequences; the plan contrast persisted at alternate cloud thresholds. The plan effect was present before the end of the common sequence. We discuss the implications for look-ahead scheduling and the distinction between announcing an extra sequence and actually planning to observe it.

I. INTRODUCTION

A cloud map is useful to an observer only if it arrives before the target is lost. All-sky optical cameras and infrared monitors supply complementary measures of cloud cover; stellar photometry supplies a check on the transmission along a particular line of sight [1–3]. Advection of cloud masks can extend the useful forecast horizon by several minutes, although cloud formation and evolution remain difficult to predict [4]. A scheduler that treats the forecast as independent of the observing programme may still be wrong if the end of the programme and the arrival of cloud are associated.

Observers often identify the next sequence as their last before closing the telescope. A concentration of interruptions at this point could arise from selection: an observer who sees a deteriorating sky is more likely to choose a final sequence. It could also arise from the declaration itself, or from the decision to stop. These explanations make different predictions when the announcement and the observing plan are assigned independently. We report a factorial experiment in which they are separated while the telescope executes the same initial sequence in every arm.

The measured quantity is the onset of cloud in a fixed sky sector, not the time at which the observer closes the dome. The distinction matters because the latter is partly a decision. The field sensor and its reduction are described in Sec. II, the assignment and analysis in Sec. III, and the results in Sec. IV. The experiment tests the immediate effect of a plan to stop, not whether an indefinitely extended programme can prevent cloud.

II. OBSERVATORY AND MEASUREMENTS

A. Telescope and cloud monitor

The study used the 1.2-m telescope at the Estmere Observatory, a remotely scheduled optical facility at 2,300 m elevation. A fixed monochrome all-sky camera acquired 20-s exposures on a 60-s cadence. We registered the images to a catalogue of bright stars, masked the dome and a 20° radius about the Moon, and retained directions above 30° altitude. The telescope field was represented by a 10°-radius sector centred on its initial pointing. A sector was ineligible if more than one quarter of its area fell inside a mask.

Cloud fraction $C(t)$ is the fraction of unmasked sector area classified as cloud in a camera frame. The classifier combines local losses of catalogue stars with spatially coherent changes in the sky background; its thresholds were fixed using commissioning images that were not used in the trial. We calibrated the camera's sky coordinates against stellar positions [5]. The sector was fixed in horizontal coordinates for the 15-min window, so telescope tracking and a change of pointing could not move the measurement aperture onto a cloud. This sacrifices exact tracking of the science field near the sector edge; none of the trial pointings left the sector in 15 min.

The primary event was the first of two consecutive frames with $C \geq 0.30$. Its time was assigned to the first frame. A trial began only if the preceding 10 frames had $C < 0.10$; the trial clock started with the announcement. A cloud event in the last minute was counted only if confirmed by the next frame. Thus each trial has 15 min of onset ascertainment and one additional minute for confirmation. The threshold refers to sector area, not optical depth. Thin uniform extinction can be missed, particularly in moonlight.

A second, independently reduced $10\,\mu\mathrm{m}$ radiometer monitored the same sector at 1-min cadence. It was not used for eligibility, randomization or the primary endpoint. Two readers, blind to assignment, inspected 200 trial sequences sampled evenly across arms and spanning lunar illumination and season. Their consensus, based on the optical sequence and radiometer trace, identified 99 events. The fixed optical rule found 93 of these and called 5 of the 101 non-events positive. These values apply to the selected sequences, not to arbitrary all-sky images.

B. Photometric check

A small co-aligned camera measured differential r-band magnitudes for bright stars in the sector. For each frame we subtracted a clear-sky reference and the median residual in three comparison sectors. The procedure reduces common extinction and zero-point drift but also removes spatially uniform cloud. It is therefore a check on localized extinction, not an independent measurement of the full cloud fraction. The median excess extinction during events identified by both the cloud monitor and the radiometer was 0.18 mag; its interquartile range was 0.11–0.29 mag. The two instruments share the atmosphere, though not their detectors or reduction pipelines.

III. ASSIGNMENT AND ANALYSIS

A. Factorial design

Between 2021 and 2025, 400 eligible nights entered the trial, at most one per night. The first eligible 15-min interval after the middle of astronomical night was selected before assignment. A candidate was excluded before assignment if the sector was masked, the monitor was unavailable, or cloud cover in the preceding 10 min failed the criterion above. Safety closures and weather limits remained in force. No assigned trial lost its cloud-monitor record during the endpoint window.

The observing plan and the spoken sentence were independently assigned in blocks of four consecutive eligible nights, with one night per cell in each block. For the **stop** plan, the operator was instructed to finish science observations after the common sequence. For the **continue** plan, the operator was instructed to carry out at least one further science sequence after it, subject to the ordinary safety rules. In both cases the common sequence consisted of three 300-s exposures in the same filter and sky sector. The programmed camera and telescope commands during the first 15 min were identical in the two plans. The additional sequence was scheduled to begin after the 15-min endpoint, and could be cancelled for weather. We analyse assigned plans, not whether the later sequence was actually completed.

At $t = 0$ the operator read either “This is the last sequence” or “The next sequence starts now.” The latter is the neutral sentence. In the continue/last cell the first sentence was deliberately false; in the stop/neutral cell it did not disclose the plan. Only the operator and scheduler knew both assignments during the trial. The camera was not connected to the scheduling computer, and the person who reviewed the monitor records received only trial identifiers. This arrangement separates a statement about stopping from the scheduled decision to stop, although it does not isolate the operator’s private expectation from the scheduler’s record of the plan.

B. Estimands and uncertainty

For plan p and sentence s, let $Y_{ps} = 1$ if cloud onset occurred within 15 min and zero otherwise, with event fraction $\hat{p}_{ps}$. Write S and C for stop and continue, and L and N for last and neutral. We estimated the marginal plan effect as the difference of the stop and continue event fractions, averaged equally over sentences:

$$\widehat{\Delta}_{P} = \frac{\hat{p}_{SL} + \hat{p}_{SN}}{2} - \frac{\hat{p}_{CL} + \hat{p}_{CN}}{2}. \tag{1}$$

The sentence effect exchanges the roles of plan and sentence in Eq. (1). We also report the difference between sentence effects in the two plan groups. With 100 nights per cell, differences and approximate 95% intervals use binomial standard errors for independent nights. Consecutive nights may share weather systems, so these intervals should be read as approximate; they do not account for correlations between nights in a block. The order of assignment within each block was randomized, and the primary endpoint and the two marginal contrasts were fixed before the assignments were opened.

We repeated the onset calculation at $C = 0.15$ and 0.50 as a sensitivity check. These are alternate definitions of the outcome, not independent replications. The event-time plot uses the same two-frame confirmation rule at 5, 10 and 15 min. We report fractions and risk differences rather than a fitted hazard ratio because the risk is appreciable over the observing interval and the operational question concerns a fixed 15-min horizon.

IV. RESULTS

A. Cloud onset

Table I gives the number of onsets in each arm. Cloud entered the sector on 91 of 200 stop nights and 36 of 200 continue nights. The estimated plan effect was 27.5 percentage points (95% interval 18.8 to 36.2). The corresponding event fractions were 45.5% and 18.0%. The effect of the sentence, averaged over plans, was 2.5 points (interval −6.6 to 11.6). The difference between the two within-plan sentence effects was 9.0 points (interval −8.4 to 26.4); the data are compatible with a modest interaction as well as none.

TABLE I. Cloud onset within 15 min by assigned plan and sentence. Each cell contains 100 nights. Fractions refer to first onset and include only events confirmed in the following frame.

Plan	Sentence	Events	Fraction (%)
Stop	Last	49	49
Stop	Neutral	42	42
Continue	Last	17	17
Continue	Neutral	19	19
All stop	Either	91	45.5
All continue	Either	36	18.0

The difference in onset fractions was already present at 5 min and widened through the endpoint (Fig. 1). In the stop arms 22 of 200 nights had an onset by 5 min, compared with 8 of 200 in the continue arms. The common sequence was still underway at that time. Changing the words while retaining the plan did not reproduce the separation between plans. In particular, a false final-sequence announcement on a continue night gave 17 onsets, compared with 19 after the neutral sentence.

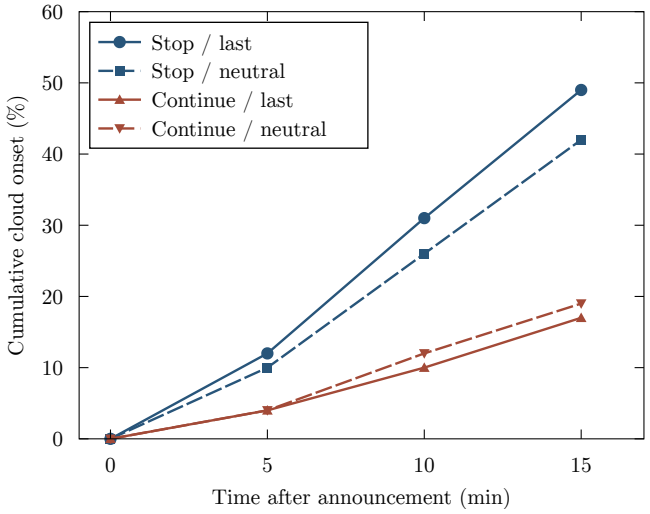


FIG. 1. Cumulative fraction of nights with confirmed cloud onset. Each curve contains 100 nights; the plan is assigned independently of the sentence. Points are at 5-min intervals and lines connect them for display.

The plan contrast persisted as the cloud-fraction threshold changed (Fig. 2). At thresholds 0.15, 0.30 and 0.50 the stop-minus-continue differences were 23.0, 27.5 and 22.0 percentage points, respectively. The less restrictive threshold includes more thin or fragmented cloud, while the strictest favours larger opaque features. Since the same sequences enter all three calculations, the similarity of the estimates does not multiply the sample size.

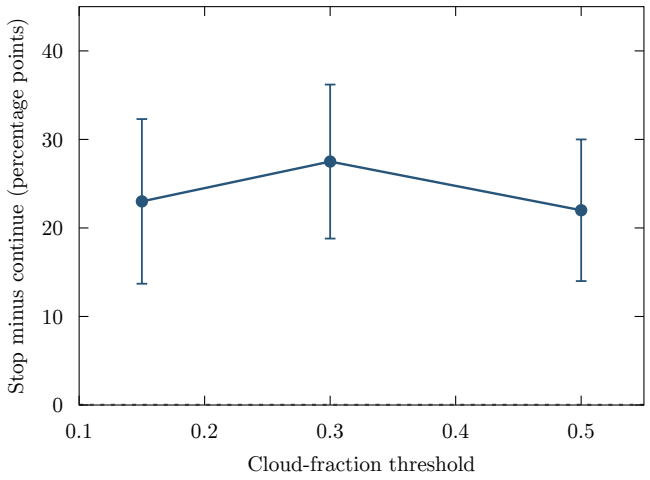


FIG. 2. Difference in 15-min onset fraction between assigned stop and continue plans as the sector cloud-fraction threshold changes. Error bars are approximate 95% binomial intervals; the middle point is the primary threshold.

B. Measurement checks and error budget

The blinded visual and radiometric review yielded a sensitivity of 93.9% and specificity of 95.0% against the consensus labels (Table II). The radiometer alone is less sensitive to some thin cloud, while the optical camera can confuse low stellar density or scattered moonlight with obscuration. Since the consensus uses both measurements, this comparison is a consistency check rather than an independent estimate of absolute truth. The reviewers saw neither the schedule nor the sentence.

TABLE II. Agreement between the fixed optical event rule and blinded consensus in 200 selected sequences. Rows are the optical rule; columns are the consensus.

Optical classification	Consensus event	Consensus no event
Event	93	5
No event	6	96

The sampling standard error of the primary risk difference is 4.4 percentage points. Moving the cloud threshold over the stated range changes the point estimate by at most 5.5 points, but this range is not a standard uncertainty and should not be added in quadrature to the sampling error. One-minute cadence and the two-frame rule limit the timing precision to about a minute; they do not change the assignment order. The trial's most consequential uncontrolled term is correlation in weather between nearby nights. The factorial assignment is balanced within four-night blocks, but a larger experiment should estimate uncertainty by block and site rather than by treating nights as independent.

V. DISCUSSION

The plan association cannot be explained by an observer choosing to stop after seeing an approaching cloud, because the plan was assigned after the clear-sky eligibility check. It is also not explained by the sentence alone in this sample. The schedule, the operator's knowledge of it and any actions correlated with that knowledge remain inseparable. The experiment does not identify a mechanism, and a single site cannot establish how far in advance a plan must be fixed or whether the result transfers to robotic operation.

The result bears on look-ahead scheduling. A forecast based only on camera images has no representation of whether the current sequence is scheduled to be the last. In the trial sample the constant forecast equal to the overall event rate, 0.3175, has a Brier score of 0.217. Using the two observed plan-specific rates gives 0.198 on the same nights. This is an in-sample calculation, not a validated forecast improvement. It illustrates the size of the information available to a scheduler; a prospective test would need forecasts fixed before the night and evaluation on later nights.

A scheduler could postpone its final sequence by planning a subsequent one. Our estimate applies only to the first such decision: in the continue arm the next sequence was a real instruction, not a placeholder, and the trial did not randomize what happened after it. In particular, the result gives no basis for assuming that adding another sequence repeatedly will keep a telescope open indefinitely. At some point the scheduler must declare the programme complete, and the 15-min cloud forecast must be recomputed.

VI. CONCLUSIONS

At Estmere, a randomly assigned plan to stop after the current sequence was followed by more cloud onsets than a plan to continue, even though the telescope and the cloud monitor followed the same commands for the first 15 min. A spoken final-sequence declaration alone produced no comparable shift. A multi-site experiment with concealed schedules and block-level uncertainty is needed before treating the plan as an input to an operational forecast.

[1] D. W. Hogg, D. P. Finkbeiner, D. J. Schlegel, and J. E. Gunn, The Astronomical Journal **122**, 2129 (2001).

[2] J. Aceituno, S. F. Sánchez, F. J. Aceituno, D. Galadí-Enríquez, J. J. Negro, R. C. Soriguer, and G. Sánchez Gómez, Publications of the Astronomical Society of the Pacific **123**, 1076 (2011).

[3] M. Mommert, The Astronomical Journal **159**, 178 (2020).

[4] S. Buntin, C. M. Copperwheat, and H. E. Jermak, RAS Techniques and Instruments **4**, rzaf034 (2025).

[5] D. Lang, D. W. Hogg, K. Mierle, M. Blanton, and S. Roweis, The Astronomical Journal **139**, 1782 (2010).