Differing Air Traffic Controller Responses to Similar Trajectory Prediction Errors

An Interrupted Time-Series Analysis of Controller Behavior

Joey Mercer (*Author*)

Airspace Operations Laboratory
NASA Ames Research Center
Mountain View, CA, USA
joey.mercer@nasa.gov

Sarah Hunt Espinosa (*Author*)

User-Centered Design & Engineering
SPAWAR Systems Center Pacific
San Diego, CA, USA
shunt@spawar.navy.mil

Nancy Bienert, Sean Laraway, San José State University / NASA Ames Research Center, Moffett Field, CA

Abstract— A Human-In-The-Loop simulation was conducted in January of 2013 in the Airspace Operations Laboratory at NASA's Ames Research Center. The simulation airspace included two en route sectors feeding the northwest corner of Atlanta's Terminal Radar Approach Control. The focus of this paper is on how uncertainties in the study's trajectory predictions impacted the controllers' ability to perform their duties. Of particular interest is how the controllers interacted with the delay information displayed in the meter list and data block while managing the arrival flows. Due to wind forecasts with 30-knot over-predictions and 30-knot under-predictions, delay value computations included errors of similar magnitude, albeit in opposite directions. However, when performing their duties in the presence of these errors, did the controllers issue clearances of similar magnitude, albeit in opposite directions? This paper describes the use of a novel technique (Interrupted Time Series) to examine the controller response data.

Keywords—ATC; Interrupted time-series

I. INTRODUCTION

The National Airspace System (NAS) forecasts continued growth in traffic demand [1], and under the plans for the Next Generation Air Transportation System (NextGen), the Federal Aviation Administration (FAA) aims to address one of the system's constraining factors, the controller's mental capacity, by increasing the use of automation aids [2, 3]. These tools depend on the predicted speed and path of an aircraft: its trajectory. Trajectory prediction capabilities are therefore a fundamental part of future Air Traffic Management (ATM) systems, and will be used by NextGen automation tools to provide advisory aids. The performance of such Decision Support Tools (DSTs), and ultimately their operational acceptance, is likely dependent on the accuracy of the underlying trajectory predictions. However, trajectory predictions, by their very nature, are not perfect: they are informed guesses.

With the predicted trajectory of an aircraft, NextGen's advanced DSTs can provide an alert: either cautionary (such as highlighting a conflict) or informational (such as suggesting a speed to meet a schedule, or displaying the predicted delay at a metering point), if needed. Uncertainties in trajectory predictions then, directly affect the DST's ability to help the controller perform their task. While it is possible that the controller could compensate for the errors in 'bad automation,' they can only do so to a limited extent. More specifically, the controller's efforts to compensate for the automation may reach a workload ceiling, at which point performance may worsen.

A. Study of Trajectory Prediction Uncertainties

A Human-In-The-Loop simulation was conducted in January of 2013 [4] in the Airspace Operations Laboratory (AOL) at NASA's Ames Research Center [5]. The simulation airspace included two en route sectors (one high-altitude and one low-altitude) feeding the northwest meter-fix of Atlanta's Terminal Radar Approach Control (TRACON), depicted in **Error! Reference source not found.** The test participants were responsible for delivering aircraft to the meter-fix within +/- 20 seconds of the scheduled time (information shown to them on their display in a meter list and in aircraft data blocks), and for providing standard separation services for all aircraft. The environment also included over-flight and departure traffic, thereby increasing the complexity of the task of providing separation.

The participants staffing the test sectors were retired air traffic controllers, none of whom were familiar with the test airspace, and had an average of 23.75 years of experience and had been retired for an average of 5.5 years. Confederate controllers, also retired, staffed Radar-Associate (D-side) positions, one for each test sector, while student/general aviation pilots staffed the confederate pseudo-pilot positions. During a one-week study, two separate simulations were conducted simultaneously and in parallel, creating two

‘worlds’. Although the approach of two worlds required twice the number of positions described above, it doubled the amount of data collected in the same amount of time.

Figure 1: Experimental sectors simulating an arrival flow into Atlanta.

The time-frame of the operations simulated in this study was anchored in a notional ‘mid-term’ NextGen environment, where better surveillance data and advanced DSTs could be expected. In performing their duties, the controllers had several automation aids available to them (see Figure 2). Meter lists contained arrival sequence and schedule information, while delay indications were also shown in the data block. A conflict probe displayed information regarding detected conflicts in a conflict list as well as in the data block. Trial-planning tools were also available, allowing controllers to manually craft trajectory changes to solve local problems. Additionally, functions were available to the controllers that allowed them to request a solution from the automation for detected problems.

Figure 2: Screenshots of mid-term operational tools.

It is important to note that the delay information displayed in the meter list and data blocks was configured with a unique behavior that could be considered a tool in its own right. Any amendments made by the controllers to an aircraft’s trajectory, either manually entered, or trial-plan-assisted, caused the ground system’s automation to immediately compute a new trajectory which incorporated the newly available information. For example, if the controller issued a speed clearance to an aircraft, when inputting the new speed as a system entry, the automation would then compute a new trajectory for the aircraft based on the new speed. This had an immediate effect on the delay information displayed in the data block and meter list, which would immediately update to reflect the new trajectory prediction.

Because all tools were trajectory-based, they were subject to the various simulated errors inherent in those trajectories, meaning that due to uncertainties in the trajectories provided, the tool information displayed to the controllers was imperfect. This highlights one of the simulation’s primary objectives: to examine at which point the automation tools would become unacceptable to the controllers and no longer support adequate system performance in terms of separation services or metering conformance.

The simulation investigated how trajectory prediction uncertainties impacted the controller’s ability to provide standard separation services and deliver aircraft on time to the meter-fix. This paper aims to understand the nature of the controllers’ response to different uncertainties simulated across similar conditions. Of interest is how the participants’ form of compensating for prediction errors affected the clearances they issued, and the exploration of new analysis methods to gain further insight into controller behavior.

B. Simulation of Trajectory Prediction Errors

Uncertainties were introduced in the form of wind forecast errors and errors in aircraft performance assumptions (e.g., climb/descent rates). A selection of different Rapid-Update Cycle (RUC) wind files created mismatches between environment and forecast wind fields. Wind forecast errors either over-predicted or under-predicted a predominant tail wind by varying amounts. A baseline condition with no wind errors was included, as well as ‘Realistic’ Root-Mean-Square (RMS) wind errors of 10 knots, meant to represent typical ‘real-world’ forecast errors. Other levels of wind error included ‘Moderate’ RMS errors of 20 knots, and ‘Large’ RMS errors of 30 knots.

The simulation also investigated errors in the underlying aircraft performance models, which were implemented such that while the ground system’s assumptions about aircraft performance remained constant, the actual descent and climb performance of individual aircraft behaved according to modified ‘scale factors.’ The scale factors were designed to impact the distance normally needed by an aircraft to descend from one altitude constraint to the next, or to climb from one altitude constraint to the next. The simulation examined a baseline condition with no aircraft performance errors, as well as two additional target levels of aircraft performance model

errors: ‘Realistic’ errors of approximately 5%, and ‘Large’ performance errors of approximately 25%.

The simulation employed two primary scenarios (scenarios A and B), designed independently, but meant to be comparable. Coupling the two scenarios with different combinations of forecast and environment (i.e., ‘truth’) winds allowed the simulation to not only examine different magnitudes of wind error, but also both directions of error bias. A positive bias (an over-prediction error) resulted when the forecast winds were stronger than the environment winds, whereas forecast winds that were weaker than environment winds resulted in a negative bias (an under-prediction error). Environment and forecast winds were paired in these two ways for each of the wind-forecast error conditions.

During the simulation, traffic scenario A was mostly paired with positive-bias wind errors, whereas traffic scenario B was mostly paired with negative-bias wind errors. The negative-bias wind errors used during trials with scenario B had the effect of presenting the controllers with seemingly smaller initial delay values. In contrast, the positive-bias wind errors used during trials with scenario A impacted the trajectory predictions such that the controllers saw seemingly larger initial delay values. For an aircraft left untouched, the delay would gradually correct towards the actual delay (i.e., the delay expected using perfect wind forecast information) as it came closer to the meter-fix. This is true in either scenario: that is, the seemingly smaller initial delay value in scenario B would gradually increase as the aircraft approached the meter-fix, and conversely, the seemingly larger initial delay value in scenario A would gradually decrease.

II. COMPARISONS MADE

The simulation results showed that scenario B was less challenging for the controllers than scenario A [4], [6]. Given the different wind-error biases associated with each scenario, the present investigation explores the relationship between the wind-error biases and the study’s findings. To achieve this, runs 11 and 12 from the study are examined because they both simulated 30-knot forecast wind errors: run 11 did so with a positive wind-error bias, while run 12 did so with a negative wind-error bias.

Central to the current analysis is the controller’s response to the trajectory prediction errors as a result of the wind-error bias. While it is true that the entire set of instructions issued by the controller to an aircraft represents that response, this analysis distinguishes the first speed clearance from the remaining clearances, thereby isolating what the authors believe best represents the controller’s initial judgement of the response needed, from later corrective actions. In this regard, although the assigned speed is informative, it comes from a limited range of flyable speeds. Therefore, the current analyses consider the relative magnitude of the issued speed change (i.e., the difference between the aircraft’s current speed and the issued speed), and when examined for all aircraft over the course of a run, provide insight in how the controller’s judgement changed over time. This approach embraces the natural learning controllers do as they issue clearances, observe their effect, and adjust accordingly for the next clearance.

Exploring assigned speeds and speed-change magnitudes together allows for a multi-dimensional analysis of the controller’s behavior not possible with just one metric.

The scenarios used in runs 11 and 12 were similar but not identical, yielding within-subject data that is not directly comparable and does not fit traditional statistical testing. Instead, an Interrupted Time-Series (ITS) analysis was used to detect any progressive changes over time for a given participant.

A. Data Analysis

This study arranged several two-phase (A-B) interrupted time-series designs in which wind-error bias changed from run 11 (the “A” phase) to run 12 (the “B” phase). To analyze our data, we used an interrupted time-series (ITS) analysis strategy based on the work of Huitema, McKean, and colleagues [7, 8, 9, 10, 11, 12]. In ITS, “time series” refers to the fact that the data are collected at regular intervals over time. The “interruption” is the onset of some event during that time series, such as a change in conditions or addition of intervention. In this study, the “interruption” was the interval between run 11 and run 12 and the subsequent change in wind-error bias.

The ITS used in this study started by fitting ordinary least-squares (OLS) regression models to the speed event and speed-change magnitude data (measured in knots). The purpose of the ITS is to determine if the interruption changes the data from one phase to the next. This involves fitting a separate regression line to each phase and testing for changes in the properties of the two regression lines. Although we can compute several measures for these regression lines, the primary measures of interest are: changes in level and slope. The meaning and interpretation of the term “slope” is the same as in other applications of OLS regression models. That is, the slope provides information on the change in Y (the dependent variable) given a one-unit change in the predictor variable X. In ITS, the X variable is time. Slope change (SC) occurs when the slope (trend) in the first phase is significantly different from that in the second phase (e.g., positive in the first, negative in the second; zero in the first, positive in the second, etc.). Level change (LC) can be computed in various ways. Following the methods described by Huitema [7], we computed level change as the difference between the value of the Y variable, as predicted from the regression model in the first phase, from that predicted from the regression model at the first time-point in the second phase. If there is no level change, then these two estimates will be the same. If there is a significant level change, this means that the interruption/intervention (in this study, the reversal of wind error), shifted the level of behavior in the second phase compared to the first.

The regression lines may have different levels, slopes (trends), or both. The data in the first phase can be used to predict data in the second phase if the interruption has no effect (called a “counterfactual”). If the interruption has no effect, then we expect the regression line in the second phase to resemble that of the first phase (i.e., show no changes in

level or slope). The extent to which the properties of the regression line in the second phase differ significantly from that of the first phase, support the possible conclusion that the interruption is at least partially responsible for the observed change (of course, alternative explanations must be ruled out to make strong claims about the effects of the interruption).

If the slopes in both phases are zero, then the level change is the same as the difference between the means of the two phases. In the presence of non-zero slopes in either or both phases, level change is not equivalent to the mean difference. The advantage of the using Huitema's ITS approach rather than simply comparing phase means is that in the presence of non-zero slope(s), the mean difference "does not convey the fact that the size of the effect (if any) is a function of the within-phase time period," "may be large even when there is no effect whatsoever," and "may not reveal an effect when one is present" [7, p. 371]. Thus, Huitema's ITS approach allows for more precise descriptions and estimates of intervention effects in a time series than are possible by simply computing differences in means.

As in other applications of OLS regression, in Huitema's ITS approach the most appropriate and parsimonious model is chosen to describe the dependent variable. The available models for this analysis are listed in Figure 3 and will be referenced numerically (1-4). Like traditional applications of OLS regression, Huitema's ITS analysis is accomplished by (1) regressing the dependent variable Y on predictor variables, and (2) evaluating the properties of the errors to assess violations of assumptions. The terms in the various models are defined as [5]:

Y_t = dependent variable score at time t
 β_0 = the predicted value of Y at time 0
 β_1 = the slope of the first (A) phase
 β_2 = the level-change regression coefficient
 β_3 = the change in slope from the first to the second phase
T_t = the value of the time variable at time t
D_t = the value of the level-change dummy variable (coded 0 for the first phase and 1 for the second phase) at time t
SC_t = the value of the slope-change variable, calculated as $[T_t - (n_1 + 1)]D$, where n_1 is the number of observations in the first phase
 Φ_1 = the lag-1 autoregressive coefficient
 ϵ_t = the error (residual) of the process at time t
 μ_t = the disturbance at time t $[Y_t - (\beta_0 + \beta_1 T_t + \beta_2 D_t + \beta_3 SC_t + \Phi_1 \epsilon_{t-1})]$

By adding β_1 and β_3 , we can obtain the value of the slope for the second phase. Disturbances and errors are assumed to be independent, normally distributed with a mean of zero, and showing a constant variance at all time points (homoscedasticity). Note the similarities to standard OLS regression.

The starting point of the analysis is to plot the raw data to examine trends and potential changes in level. The next step involves fitting Model 1 and Model 2 to determine which model provides the better description of the data. The choice between models is based on whether slope and slope change are non-zero (i.e., $\beta_1 = \beta_3 = 0$). If they are both non-zero, then Model 1 is preferred because it incorporates information about slope and slope change. If they are both zero, then Model 2 is preferred because the slope and slope change parameters are unnecessary. Huitema [7] described a model comparison test to determine which model should be chosen. The null hypothesis for this test states that the slope and slope change coefficients in Model 1 are equal to zero ($H_0: \beta_1 = \beta_3 = 0$). Rejecting this hypothesis indicates that the two coefficients are necessary to describe the data adequately; thus, Model 1 should be chosen. If this hypothesis is not rejected, then Model 2 should be chosen because the slope and slope change coefficients are not necessary to describe the data adequately; the simpler Model 2 is sufficient. In this case, Model 2 essentially tests the difference in the two phase means. The model comparison test statistic is computed as:

$$F = \frac{(SS_{Reg\ Model\ 1} - SS_{Reg\ Model\ 2})/2}{MS_{Residual\ Model\ 1}}$$

Where:

F = test statistic
SSReg Model 1 = regression sum of squares from Model 1
SSReg Model 2 = regression sum of squares from Model 2
MSRes Model 1 = residual mean square from Model 1

This test statistic is compared to a critical value having $df = 2$, $N - 4$, where N is the total number of observations. Huitema [7] recommends using a liberal alpha level (e.g., 0.10) for this test. Once the appropriate model has been chosen, its errors should be evaluated to determine if they conform to the assumptions described above. The methods for doing so are the same as with other applications of OLS regression models, with one important exception: auto-correlation. Because the data are collected over time from the same individual or group of individuals, it is possible that the errors of the regression model show serial dependence or auto-correlation. That is, an error at one point in time provides information on the value of errors at other points in time. There are several ways to determine if significant auto-correlation is present in a time series, but a full discussion is beyond the scope of this paper. The interested reader is referred to [7, pp. 378-382, 9, 10].

Should the errors from either Model 1 or Model 2 show significant auto-correlation, then Models 3 or 4 are used. Model 3 and Model 1 are the same except that Model 3 adds a regression coefficient to account for auto-correlation among the errors (i.e., the "auto-regressive coefficient"). So, if Model 1 was chosen and its errors show a significant lag-1 auto-correlation coefficient, then Model 3 should be used.

Similarly, Model 2 and Model 4 are the same, but Model 4 includes an auto-regressive coefficient. The software to fit Models 3 and 4 is based on the procedure described in [12] and can be found at this website:

<http://www.stat.wmich.edu/slab/Software/Timeseries.html>.

Once the appropriate model has been chosen, effect size measures can be computed to describe the magnitude of the observed changes in the dependent variable. Huitema [7] defined one such measure, the standardized level change (SLC). The SLC expresses the size of the level change in standard deviation units. This measure is computed as:

$$SLC = \frac{b_{LC}}{\sqrt{MS_{Res}}}$$

Where:

bLC = the estimate of the level change for the final model chosen (b2 for Models 1 and 3; b1 for Models 2 and 4)

MSRes = residual mean square from the final model

III. RESULTS

A. World 1, sector 5

Table 1: World 1, Sector 5's results.

World 1, Sector 5	Speed Events	Speed Change Magnitude
Model	1	2
Level Change (D)	no	yes
Baseline Slope [Run 11] (T)	yes	--
Slope Change (SC)	yes	--
Autocorrelated?	no	no

W1S5 did not show a significant level change from runs 11 to 12 in their raw speed event data $\beta = -3.58$, $t(112) = .40$, $p = .69$, SLC= .15, however the level change for speed-change magnitude was significantly different $\beta = -1.04$, $t(114) = -5.57$, $p < .001$, SLC= -1.04. As seen in Figure 3, there was a significant slope in the baseline of speed events $\beta = .47$, $t(112) = 2.63$, $p = .001$, SLC= .15 and significant speed event slope level change $\beta = -.58$, $t(112) = -2.15$, $p = 0.03$, SLC= 0.15 from runs 11 to 12.

Figure 3: World 1, Sector 5's graphical display of the chronological data.

B. World 2, sector 5

Table 2: World 2, Sector 5's results.

World 2, Sector 5	Speed Events	Speed Change Magnitude
Model	2	1
Level Change (D)	yes	yes
Baseline Slope [Run 11] (T)	--	no
Slope Change (SC)	--	yes
Autocorrelated?	no	no

W2S5 had a statistically significant level change for both speed event data $\beta = 12.39$, $t(60) = 2.083$, $p = 0.04$, SLC= 0.53 and speed-change magnitude $\beta = -40.86$, $t(58) = -4.55$, $p < .001$, SLC= -2.32, which can be seen in Figure 4. Model 2 fully explained the speed event data, removing slope. However, the speed-change magnitude data required model 1, and while there was no statistically significant slope in run 11, there was a statistically significant slope change from runs 11 to 12 $\beta = 1.46$, $t(58) = 2.77$, $p = 0.008$, SLC= -2.32

Figure 4: World 2, Sector 5's graphical display of the chronological data.

C. World 1, sector 6

Table 3: World 1, Sector 6's results.

World 1, Sector 6	Speed Events	Speed Change Magnitude
Model	2	2
Level Change (D)	yes	yes
Baseline Slope [Run 11] (T)	--	--
Slope Change (SC)	--	--
Autocorrelated?	no	no

W1S6 used model 2 to explain both data sets, with level change the only variable of interest (see Figure 5). For both speed events $\beta = 47.89$, $t(55) = 9.37$, $p < .001$, SLC= 2.60 and speed-change magnitude $\beta = 36.91$, $t(55) = 5.27$, $p < .001$, SLC= 1.44 there was a statistically significant level change between runs 11 and 12.

Figure 5: World 1, Sector 6's graphical display of the chronological data.

D. World 2, sector 6

Table 5: World 2, Sector 6's results.

World 2, Sector 6	Speed Events	Speed Change Magnitude
Model	3	3
Level Change (D)	yes	no
Baseline Slope [Run 11] (T)	no	yes
Slope Change (SC)	no	yes
Autocorrelated?	yes	yes

W2S6 had the only instance of auto-regressive errors, and with the correction, both data sets used model 3. Level change was statistically significant in speed event $\beta = 34.09$, $t(59) = 2.02$, $p = 0.05$, SLC = 1.87 but not speed-change magnitude $\beta = 1.29$, $t(59) = 0.01$, $p = 0.92$, SLC = 0.06. Speed event data did not have a statistically significant baseline slope $\beta = 0.8$, $t(59) = 1.13$, $p = 0.26$, SLC = 1.87 or slope change $\beta = -1.35$, $t(59) = -1.06$, $p = 0.29$, SLC = 1.87. However, speed-change magnitude had both a statistically significant baseline slope $\beta = 1.60$, $t(59) = 3.6$, $p < 0.001$, SLC = 0.06 and slope change $\beta = -1.63$, $t(59) = -2.04$, $p = 0.05$, SLC = 0.06 (see Figure 6).

Table 4: Statistical summary of the interrupted time series analysis.

	Significance				Effect Size- R ²					Standardized Level Change Statistic		
Measure	Beta	t	df	p	R ²	F	df1	df2	p	B _{lc}	MSresidual	SLC Statistic
World 1, Sector 5: Model 1: Speed Events												
Level Change	3.58	0.40	112.00	0.69	na	na	na	na	na	3.5819	581.2249	0.1486
Baseline Slope	0.47	2.63	112.00	0.01*								
Slope Change	-0.58	-2.15	112.00	0.034*								
World 1, Sector 5: Model 2: Speed Change Magnitude												
Level Change	-1.0351	-5.571	114	0.00**	0.2400	31.036	1	114	0.00	-15.6670	229.0731	-1.0351
Baseline Slope	na	na	na	na								
Slope Change	na	na	na	na								
World 2, Sector 5: Model 2: Speed Events												
Level Change	12.3915	2.083	60	0.041*	0.0670	4.341	1	60	0.041	12.3915	539.1922	0.5336
Baseline Slope	na	na	na	na								
Slope Change	na	na	na	na								
World 2, Sector 5: Model 1: Speed Change Magnitude												
Level Change	-40.8620	-4.552	58	0.00**	na	na	na	na	na	-40.8620	310.5021	-2.3189
Baseline Slope	0.17	0.577	58	0.566								
Slope Change	1.455	2.766	58	0.008*								
World 1, Sector 6: Model 2: Speed Events												
Level Change	47.8920	9.372	55	0.00**	0.6150	87.826	1	55.00	0.00	47.8920	339.0470	2.6010
Baseline Slope	na	na	na	na								
Slope Change	na	na	na	na								
World 1, Sector 6: Model 2: Speed Change Magnitude												
Level Change	36.9135	5.274	55	0.00**	0.3360	27.811	1	55.00	0.00	36.9135	636.0740	1.4636
Baseline Slope	na	na	na	na								
Slope Change	na	na	na	na								
World 2, Sector 6: Model 3: Speed Events												
Level Change	34.0865	2.023	59	0.0476*	na	na	na	na	na	34.0865	331.6938	1.8716
Baseline Slope	0.8017	1.127	59	0.2643								
Slope Change	-1.3509	-1.064	59	0.2917								
World 2, Sector 6: Model 3: Speed Change Magnitude												
Level Change	1.2879	0.098	59	0.9223	na	na	na	na	na	1.2879	426.9135	0.0623
Baseline Slope	1.6001	3.596	59	0.0007*								
Slope Change	-1.6282	-2.036	59	0.0462*

Figure 6: World 2, Sector 6's graphical display of the chronological data.

IV. DISCUSSION

Discussion of this data centers on the individual air traffic controllers' responses to the change in wind forecast error. The primary goal was to investigate whether an interrupted time-series analysis would prove beneficial in understanding air traffic controller behavior. While this analysis does not lend to a direct comparison between air traffic controllers, examining the individual behavior does highlight potential trends. Of the data presented, speed-change magnitude is possibly more interesting, as this highlights the controller's clearances in relation to their previous decisions, allowing a visual assessment of their strategy as they adjusted their speed clearances over time. Both data types together, the raw data (speed events), and speed-change magnitude, provide a more complete picture of the controller's behavior. The analyses presented here also provide the opportunities for discussion and refinement of its application to the air traffic control domain.

The concept of level change, and whether or not their initial clearance after the interval was predictable based on their clearance strategy for run 11 (analogous to a previous shift), reveals if the controllers carried over their strategy despite the interval. World 1, Sector 5's (W1S5) level change provided insight into a controller's response to the interval when both speed events and speed change magnitude are looked at in conjunction. The first clearance given by W1S5 after the interval fell in the predicted range (no level change), suggesting no change in strategy. However, their change in magnitude from the previous aircraft was significantly different from predicted – notably changing from +11 to -9, instead of incrementing in a positive direction like the end of run 11, suggesting the controller reversed the direction of their strategy, a pattern that continued for multiple following clearances. This behavior falls in line with the bias seen in the delay information displayed to the controllers: the underlying trajectory predictions assumed the aircraft were traveling more quickly in run 11, and less quickly in run 12. W1S5 adjusted their initial speed clearances downward as they explored the automation's error and sector conditions. W1S6, W2S5, and W2S6 all had significant level changes between runs 11 and 12, suggesting that they adjusted their strategy post-interval. This could indicate a 'reset' of their expectations about tool

accuracy when approaching a new condition. Slope, which was not present in all cases, indicates a controller adjusting their clearances as they received more data. W1S6 and W2S6 both showed significant changes in their slope for speed-change magnitude during run 12, potentially indicating they were adjusting their strategy throughout the run. A visual assessment of the data shows that some controllers incremented their clearances, progressively increasing or decreasing as they learned, while some made more drastic choices.

Additional visual inspection of the data begs further analysis of the variance between clearances, something outside of the bounds of the current analysis. During the first 3 clearances of run 11, W1S5 made large jumps, a pattern also observed in W1S6 and in W1S5 (though only for speed-change magnitude). This seems to be similar to an archer gauging the distance to a target, where the first two arrows under and overshoot, with the third arrow reaching its destination. Another analogy could be that of the goldilocks principle – the bed is too long, too short, and then 'just right'. To some extent, the controllers maintained this pattern into run 12, after the interval, revealing more about the methods controllers use to gauge conditions in their sectors than about their use of a particular tool. They may be running a goldilocks test on both the sector conditions and the behavior of their system's automation. Visual inspection, combined with the presence of some significant slope changes also suggest run 12 (the negative forecast bias), may have required more refinement by the controllers to hit the 'just right' initial speed clearance. Again, further analysis of the data's variability will prove beneficial in this respect.

V. CONCLUSION

The interrupted time-series analysis (ITS), normally an experimental tool for the medical field, has potential to benefit the air traffic control research community. ITS can identify aspects of a data set otherwise difficult to see with traditional comparisons of means. As noted in the discussion, level change and slope change, while useful, do not tell the full story. Using an ITS analysis revealed that an additional analysis of the progressive behavior's variance could gather actionable insight for tool design or training. Additional thought should be given to an analysis of learning, and the impact this ITS perspective could have on a human-automation teaming.

Nevertheless, the current analysis, if only by visual inspection, does lend credence to the conclusion that the controllers responded differently to the varying forecast conditions between the two runs. The speed-change magnitude data especially highlights the slope and variability of the decision-making process as they assessed the wind bias post-interval. This initial ITS analysis indicates the following: First, the controllers were responding differently between the two conditions to compensate for the forecast error, and with individual methods. Secondly, controllers are likely to have strategies for assessing the conditions of their sector and the behavior of their system's automation independent of the current state of the data. Lastly, an ITS analysis provides a

unique look into air traffic control data that opens new avenues for further analysis and thought.

VI. REFERENCES

- [1] Federal Aviation Administration (2013). *FAA Aerospace Forecast, Fiscal Years 2013-2033*. FAA, Washington, D.C.
- [2] Federal Aviation Administration (2013). *NextGen Implementation Plan 2013*. FAA, Washington, D.C.
- [3] Joint Planning and Development Office (2010). *Concept of Operations for the Next Generation Air Transportation System*, Version 3.2. JPDO, Washington, D.C.
- [4] Mercer, J., N. Bienert, A. Gomez, S. Hunt, J. Kraut, L. Martin, et al. (2013). *The Impact of Trajectory Prediction Uncertainty on Air Traffic Controller Performance and Acceptability*. AIAA AIAA-2013-4320, American Institute of Aeronautics and Astronautics, Reston, VA.
- [5] Prevôt, T., N. Smith, E. Palmer, T. Callantine, P. Lee, J. Mercer, et al. (2014). *An Overview of Current Capabilities and Research Activities in the Airspace Operations Laboratory at NASA Ames Research Center*. AIAA-2014-2860, American Institute of Aeronautics and Astronautics, Reston, VA.
- [6] Bienert, N., J. Mercer, J. Homola, S. Morey, and T. Prevot (2014). Case Study: Influences of Uncertainties and Traffic Scenario Difficulties in a Human-in-the-Loop Simulation. *Proceedings of the 6th International Conference on Research in Air Transportation (ICRAT)*, Istanbul, Turkey.
- [7] Huitema, B. E. (2011). *The Analysis of covariance and alternatives: Statistical methods for experiments, quasi-experiments, and single-case studies*. Hoboken, NJ: John Wiley & Sons, Inc.
- [8] Huitema, B. E., and Laraway, S. (2006). Single-subject designs. In N. J. Salkind (Ed.), *Encyclopedia of Measurement and Statistics*. (pp. 904-907). Thousand Oaks, CA: Sage Publications.
- [9] Huitema, B. E., and McKean, J. W. (2000). Design specification issues in time-series intervention models. *Educational and Psychological Measurement*, 60, 38-58.
- [10] Huitema, B. E., and McKean, J. W. (1998). Irrelevant autocorrelation in least squares intervention models. *Psychological Methods*, 3, 104-116.
- [11] Huitema, B. E., McKean, J. W., & Laraway, S. (2007). Time-series intervention analysis using ITSACORR: Fatal Flaws. *Journal of Modern Applied Statistical Methods*, 6, 367-379.
- [12] McKnight, S. D., McKean, J. W., and Huitema, B. E. (2000). A double bootstrap method to analyze linear models with autoregressive error terms. *Psychological Methods*, 5, 87-101.

VII. ACKNOWLEDGEMENTS

This work was conducted as an extension of the Functional Allocation and Separation Assurance work conducted in NASA Ames' Airspace Operations Laboratory. The NASA Airspace System Program, Concept and Technology Development Project sponsored the experiment reported in this paper. The authors wish to acknowledge the efforts of the Airspace Operations Laboratory experimental and software teams for their contributions and support. The authors would also like to thank Brad Huitema and Joe McKean for their assistance with the data analysis.