**The Autonomous Decision System Choice**

Robin Dillon[1], Peter Madsen[2], Bart Roets[3], Taylan Topcu[4], Kostas Triantis[4]

*Abstract*

Research presented previously at the Workshop on the Economics of Information Security (WEIS) has explored information security and privacy issues focusing on understanding the incentives and tradeoff decisions of people interacting with systems. This paper extends this research approach to a new domain, the tradeoff decisions regarding the use (or not) of autonomous decision-making systems (ADS). An ADS, in practice, often substitutes for a human decision maker such as in the case of an autonomous vehicle, or assists humans in complex tasks that are difficult to fully automate. ADS technology is increasingly playing a critical role in many organizational processes especially for safety-critical systems, such as controlling vehicles in transportation systems. This paper analyzes data from traffic control centers for INFRABEL, the Belgian National Railroad Company. At INFRABEL, operational service delivery decisions are made in real-time by human decision-makers, Traffic Controllers (TCs), each of whom are paired with a collaborative ADS. The ADS can set routes and open signals automatically when a train approaches or departs a station. Alternatively, the TC may decide to switch off the ADS for a certain train or sections of the railroad network and make decisions manually. This research explores which system factors are meaningful predictors of the use of the ADS by a TC in the control center versus the TC making manual decisions. The results provide insights that are applicable to the broader understanding of people's decision choices when actions can be manual or autonomous.

[1] Georgetown University
[2] Brigham Young University
[3] INFRABEL
[4] Virginia Tech

## Introduction

An ADS may be distinguished from traditional process automation in that the ADS involves higher-order, judgement-based decision-making commonly with significant financial, operational, or safety implications (McKinsey, 2019). Examples of ADS include self-driving vehicles, automated loan approval, autonomous candidate selection, advanced robotics, e-negotiation, autonomous infrastructure systems, autonomous drones, and autonomous weapon systems. Recent surveys find that ADS is one of the most commonly adopted forms of artificial intelligence with roughly half of large private-sector organizations reporting at least one form of ADS implementation (Accenture, 2019; McKinsey, 2019). A key feature of many ADS implementations is that the systems make decisions but with significant interactions with humans who hold responsibility to monitor the decisions made by the ADS (Bellet et al., 2019; Endsley, 2017).

Research on ADS is exploring human-ADS interactions and the challenges posed for ADS adoption, use, and effectiveness, particularly with regard to autonomous vehicles (Gkartzonikasa & Gkritza, 2019). Today's autonomous systems are leveraging higher levels of computational intelligence and learning algorithms than the automation of the 1980's-1990's. While building more intelligence into the systems, most systems still operate well for a range of situations that they are designed to address and need human intervention to handle non-routine, unique situations not built into their design (Endsley, 2017; Woods & Cook, 2006). We seek to understand which factors will predict when traffic controllers will intervene manually and when they will defer to the ADS by observing their behavior through data collected describing human-ADS actions. These factors will help us understand the tradeoffs that traffic controllers are making in their control decisions.

## Data

We examine the choices of human controllers to delegate a portion of their work to an ADS using a unique data set collected from a large European rail system operator, INFRABEL. INFRABEL is a government-owned corporation that builds, maintains, and operates the rail network in the country of Belgium, managing the development, repair, operation, and traffic control of the extensive and dense rail network in that country. INFRABEL's network encompasses about 3,600 kilometers of railway lines, 4,000 track-switching points and 10,000 rail signals. INFRABEL employs about 12,000 employees who build, maintain, and operate this network. INFRABEL's traffic control functions are carried out at several traffic control centers (TCCs) that collectively direct the movements of more than 4,000 trains per day. Responsibility for managing this traffic falls on traffic controllers who are assigned to oversee and direct traffic in a given portion of the rail network during a work shift (usually about 8 hours). Individual controllers cycle into and out of TCCs that run 24 hours per day 7 days per week. During a shift, controllers occupy a workstation that contains multiple computer monitors to display data concerning the portion of the network that is a given controllers responsibility as well as some information about surrounding regions of the network. Controllers direct trains by entering commands into a computer terminal to open and close signals (the equivalent of traffic lights for road traffic), switch tracks, reroute trains, and give directions to train drivers.

In recent years, INFRABEL has developed and has rolled out an ADS capable of automating the opening and closing of traffic signals—one of the biggest portions of controllers' jobs. Importantly for our purposes in this study, as INFRABEL has implemented the ADS, it has made the use of the ADS completely voluntary thus far where it is at the discretion of the traffic controller. Controllers may choose to turn over any portion of their signals to be managed by the ADS for any amount of time that they choose. They are also free to maintain full manual control of all signals if they prefer. Thus, controllers may, at their discretion, delegate to the ADS the responsibility over certain signals or certain trains for any given period of time. With this level of discretion given controllers regarding ADS use, in practice, delegation to the ADS varies significantly among controllers and over time. The goal of this study is to understand the conditions that influence controllers' decisions to delegate to the ADS or to retain manual decision making control realizing that there will be tradeoffs between delegating and retaining control.

The data used in this study cover observations of each individual work-station in each of seven INFRABEL traffic control centers for an anonymized month in 2018. Each observation in the data set covers a one-hour block of activity at a given work station, including information on ADS use, traffic conditions, controller decisions, other controller activities aside from managing signals, and controller characteristics for that hour on that workstation. The unit of analysis was the work-station-hour. Across all seven traffic control centers, an average of roughly 31 workstations were in use during a particular hour. In total, the sample contained data on 21,930 work-station-hours for the whole month.

## Variables and Analysis

**Dependent variable.** In addition to managing traffic signals, controllers at INFRABEL have two additional primary responsibilities: 1) issuing orders to route, reroute, split, or merge trains, and 2) forecasting the future state of the rail network to proactively avoid future problems by altering train routes or ordering. While traffic signal movements could be assigned to the ADS, the other two activities must be completed manually. Given that our interest in this study is how much authority human operators choose to delegate to an ADS, our dependent variable was constructed only with reference to signal movements. Specifically, the dependent variable, denoted *fraction ADS*, was the fraction of all signal movements during the workstation-hour that were performed by the ADS system rather than by the operator. This variable ranged from zero (when the ADS was not used at all during the hour) to one (when all signal movements for the hour were carried out by the ADS) with an average of 0.66. Thus, on average, controllers delegated about two thirds of signal movements to the ADS during the sampling period.

**Independent variables.** Our independent variables are key factors that may have influenced an operator's decision to delegate to the ADS or not. One such factor is signal workload. We included two independent variables to represent this factor. First, *signal workload*, was a count of the number of signal movements carried out at a workstation in a particular hour. This variable ranged from 0 to 193, with an average of 44. Second, we included a measure of the number of delays present in the portion of the network that a given workstation was responsible for during a given hour as a second indicator of workload because as delays increase, the signal movement requirements go up. This variable, denoted *traffic delay*, could be negative (when trains were ahead of schedule) or positive (when they were behind schedule) and was measured as the total time that trains were behind schedule, in hours (min: -6.9, max: 7.2, avg: 0.1).

Another factor that we investigated was the difficulty of the work faced by a controller during a given hour, and we included three independent variables as indicators of work difficulty. The first, *traffic complexity*, was an internal INFRABEL measure of how non-routine the operations in the area were for the hour (min: 0, max: 72, avg: 0.9). A second variable, *traffic density*, was an indicator of how dense the rail traffic was in the area of the workstation's responsibility based on the number of train movements per geographic area (min: 0, max: 11, avg: 0.8). The third work difficulty variable, *coordination load*, was a measure of the amount of coordination between controllers that was required during an hour to account for traffic moving from one area of responsibility to another. This variable was measured by the number of phone calls made during the hour between an operator and other operators in the TCC. The variable ranged from zero to 27 calls per hour, with an average of 1.4 per hour.

A third factor of interest in the study was the characteristics of the controllers themselves. We included two independent variables dealing with controller characteristics. The first, *time at station*, was a count of the number of consecutive hours prior to a given hour that a controller was working at the same workstation. We were unable to measure how far into a work shift a controller was before starting work at a given workstation (controllers can move from one station to another during a shift). Thus, time at station is our best indicator of how long a controller has been working and may account for the possibility that operators become more fatigued as time at the workstation increased. Time at workstation ranged from 0 to 16 hours, with an average of 2.5 hours. A second controller characteristic that was included as an independent variable was operator tenure at INFRABEL (measured in years). This variable accounts for the possibility that more experienced controllers choose to delegate less to the ADS than less tenured controllers. Operator tenure ranged from 0.6 years at INFRABEL to 42 years, with a mean of about 21 years.

In addition to the factors discussed above, we were interested in whether characteristics of the previous hour of work at a workstation continued to influence ADS use during the current hour. Thus, we included lagged values for our indicators of workload and traffic difficulty (signal workload, traffic delay, traffic complexity, traffic density, and coordination load) in the models. Additionally, we included one more independent variable accounting for the prior hour of work. This variable, denoted *errors made*, accounted for the number of errors committed by the operator during the previous hour of work. All of the work activities performed by INFRABEL controllers were carried out through commands entered into INFRABEL's control operation computer system. This system recorded all commands made by each work station in each traffic control center during the study period, including incorrect or improperly applied commands. Errors made is a count of the number of incorrect or inappropriate commands entered by a controller during the previous hour. This variable accounts for the possibility that when errors are made, controllers may choose to delegate more or less to the ADS. In our sample, *errors made* ranged from zero to ten per hour with an average of 0.13 per hour, indicating that errors were relatively rare.

**Control variables**. We also included in the analysis a set of control variables to control for the effects of other factors that might play a role in the choice to delegate to the ADS. First, we included two control variables that account for work performed by controllers other than signal movement during a given hour. In addition to signal movement, operators spend time on forecasting future rail traffic conditions to plan for future operations, so our control variable, *forecasting work,* was the amount of time, in minutes, that the operator spent during the hour forecasting future traffic conditions. Another activity that controllers performed in addition to controlling signal movement was to send commands to route and reroute trains. Thus, a second variable, *routing work,* which was an indicator of the number of routing orders issued by an operator during a given hour was also included in the analysis. A third control variable accounted for the prevalence of delegation to the ADS in the TCC in which a workstation was located. This variable, denoted *ADS use in unit*, was the fraction of all signal movements that were carried out at other workstations in the TCC during the hour. This variable may account for rail network conditions not controlled for by other variables that influence ADS use, as well as for social pressures in a TCC to use or not use the ADS.

Traffic levels and conditions on the INFRABEL network varied significantly by time of day. To account for such time of day effects, we included fixed hour-of-day effects. These fixed hour effects controlled for all factors that varied by hour of the day that may have influenced error rates, including both aspects of traffic conditions as well as potentially circadian rhythm effects on error rates. Similarly, rail traffic varied somewhat by day of the week. Thus, we also included fixed day of week effects in all models. Finally, as different traffic control centers managed different parts of the rail network with potentially different conditions that could have influenced error rates, we included fixed traffic control center unit effects in all models.

**Analysis**. As the dependent variable in the analysis was continuous, the analysis was conducted using fixed-effect linear regression. As noted above, fixed TCC effects were used to account for time invariant unit-level factors that may have influenced the choice to delegate to the TCC.

**Results**

The results of the analysis are reported in Table 1. Model 1 includes only the control variables. As can be seen, in Model 1 *forecasting work* has a positive impact on ADS use, while *routing work* has a negative effect. However, in later models, the effect of forecasting work becomes non-significant, while that of routing work remains significant. Thus, the amount of forecasting work done by a controller appears not to play a role in ADS use, while the amount of routing work done seems to reduce a controller's reliance on the ADS system. *ADS use in unit* has a positive and significant effect on ADS use, suggesting that controllers are more reliant on the ADS when other controllers in the same TCC use the ADS more.

Model 2 introduces the two independent variables tracking the amount of signal work the needs to be performed in an area. Both *signal workload* and *traffic delay* have positive and significant impacts on ADS use. This suggests that as an operator's workload goes up, their delegation to the ADS also increases. Model 3 includes the three variables that accounted for different aspects of the difficulty of the traffic that a controller was managing. All three of these variables, *traffic complexity, traffic density, and coordination load*, have negative and significant effects on ADS use. These findings indicate that as the difficulty and complexity of work to be performed increases, controllers choose to rely less on the ADS.

Model 4 introduces the two independent variables accounting for controller characteristics. *Time at workstation* has a positive and significant impact on ADS use, suggesting that as the amount of time that a controller spends at a workstation increases, ADS use also increases. This may indicate that as controllers become more fatigued, they delegate more to the ADS. *Operator tenure* has a negative and significant effect on ADS use. This finding indicates that more tenured operators may trust the ADS technology less than do less-tenured operators.

Finally, Model 5 introduces the lagged variables. We expected that traffic conditions may carry over from one hour to the next in operators' minds such that conditions in the previous hour drive ADS use in the current hour even once current traffic conditions are controlled. Indeed, the lagged workload variables both have significant, positive effects on ADS use, suggesting that having had a high workload in the past increases ADS use above and beyond the effect of current workload. Similarly, traffic complexity in the prior hour has a negative and significant effect on ADS use, although the other two lagged work difficulty variables are non-significant. Thus, there is some evidence that having dealt with complex work previously decreases the tendency to delegate to the ADS even once the complex of current work is controlled. Additionally, *errors made prior hour* has a negative and significant effect on ADS use, indicating that when an operator committed an error in the past, they become less reliant on the ADS for current work. This finding may suggest that operators blame the complexity of interacting with the ADS for errors they commit and reduce their trust in the ADS as a result. Taken together, the results of Model 5 indicate that not only current conditions, but also past conditions play a role in delegation of an ADS system.

**Discussion**

In a complex system such as a traffic control center, the TC's are constantly making tradeoff decisions regarding delegating control to the ADS or maintaining control. We summarize the factors influencing the decisions where the results from this study are consistently in the direction that would be predicted.

TC's were using the ADS more when other controllers in the same TCC are (*ADS use in unit*), when the operator's workload goes up (*signal workload* and *traffic delay*), when the operator's fatigue may be increasing (*time at workstation*), and when the workload in the past has been high (*signal workload past hour*). That operators would be influenced by other controllers is consistent with social norm theories (Cialdini, 2007, Kahneman and Miller, 1986) that have shown that people look to the actions and behaviors of others to determine their own and will conform their behavior to that of others in a group (Asch, 1956). That operators would attempt to reduce cognitive resources when their workloads are high or they are more fatigued is consistent with past cognitive load research (Shiv & Fedorikhin, 2014). Relying on the ADS when workloads are higher or fatigue is greater leaves human decision makers with more cognitive resources to handle their remaining tasks.

TC's were using the ADS less when there was greater operator's tenure, the difficulty of the task was higher as measured by *traffic complexity*, *traffic density*, and *coordination load*, the complexity in the prior hour was higher, and when errors were made in the prior hour. The finding that the operator's tenure decreases the use of ADS is related to the use of automated systems when expertise is a factor. At this time, there is only limited research that seeks to understand the role of expertise in decision making regarding the choice to use an ADS (Logg et al, 2019). These results suggest that operators with higher tenure will have more experience managing the system prior to the availability of the ADS, and may rely more on prior work routines that do not involve the ADS. Thus, here expertise is leading to a decrease in use of ADS, but more research needs to be done to understand the complex role of expertise.

Our finding that difficulty of the task and situational complexity decrease ADS use is consistent with the results of studies of other automated decision systems in particular financial robo-advisors (Fisch et al., 2018, Rühr et al, 2019). For financial advising, human advisors can offer a more personalized approach on a broader range of topics while robo-advisors are generally limited to standardized results. Therefore, for complex situations, human advisors are still necessary and the current trend in robo-advising is a hybrid approach where hybrids can charge lower fees than traditional advisors by automating part of the investment process but still providing the option of talking with a financial advisor. This hybrid approach is analogous to the voluntary use of ADS provided by INFRABEL to their traffic controllers, and best practice from robo-advising shows that the more complex situations should still be handled by the human advisors (Fisch et al., 2018). Therefore, it would make sense that when the traffic control system is in a state of higher situational complexity, ADS use would decrease.

That errors made in the prior hour, our final factor, would decrease the use of ADS is also consistent with other automation research that examines the challenge of the loss of situational awareness when control is passed between automation and human operators (Bainbridge, 1982). Further research is needed to understand how a traffic controller interprets an error and under what conditions does the controller attribute some responsibility for the error to the ADS system, but a significant challenge when automated systems control more aspects of decision making is that human controllers lose situational awareness. Therefore, with less situational awareness, it can be harder for a controller to recover from some errors. Again, we have no data to understand how the controller is attributing blame for an error occurring but concern for a loss of situational awareness could explain why when an operator committed an error in the past, they become less reliant on the ADS for current work.

**Conclusions**

Choosing to use or not use the ADS is a tradeoff influenced by many factors. The main tradeoff that we saw was that controllers were more inclined to manually control the system when the situation was more complex and more likely to defer to the ADS system when they are tired. This is an interesting tradeoff because the most critical safety situations will occur when both of these factors are present, i.e., the situation is complex and the controller is tired. We have no opinion on what is the optimal level the TC should defer to the ADS at this time, but further research can explore the impacts of the manual-ADS tradeoff on safety. If more use of the ADS is recommended at critical times, additional training may be needed to build confidence in controllers regarding the use of the ADS in complex situations. By understanding the influence of the different factors and the tradeoffs that controllers are making in their decision choices, decision makers can continue to develop training of controllers and capabilities of the ADS to optimize the use of the system to maximize efficiency and safety.

**Table 1. Fixed-Effects Linear Regression Models of Automation Use**

| | *Model 1* | | *Model 2* | | *Model 3* | | *Model 4* | | *Model 5* | |
|---|---|---|---|---|---|---|---|---|---|---|
| *Independent Variables* | | | | | | | | | | |
| Signal Workload (# movements) | | | 0.004 | *** | 0.004 | *** | 0.004 | *** | 0.003 | *** |
| | | | (0.000) | | (0.000) | | (0.000) | | (0.000) | |
| Traffic Delay | | | 0.006 | *** | 0.008 | *** | 0.008 | *** | 0.007 | *** |
| | | | (0.001) | | (0.001) | | (0.001) | | (0.001) | |
| Traffic Complexity | | | | | -0.032 | *** | -0.032 | *** | -0.027 | *** |
| | | | | | (0.001) | | (0.001) | | (0.001) | |
| Traffic Density | | | | | -0.033 | *** | -0.032 | *** | -0.030 | *** |
| | | | | | (0.002) | | (0.002) | | (0.002) | |
| Coordination Load (# calls) | | | | | -0.011 | *** | -0.012 | *** | -0.011 | *** |
| | | | | | (0.001) | | (0.001) | | (0.001) | |
| Time at Workstation (hours) | | | | | | | 0.003 | *** | 0.002 | ** |
| | | | | | | | (0.001) | | (0.001) | |
| Operator Tenure (years) | | | | | | | -0.001 | *** | -0.001 | *** |
| | | | | | | | (0.000) | | (0.000) | |
| Signal Workload, Prior Hour | | | | | | | | | 0.001 | *** |
| | | | | | | | | | (0.000) | |
| Traffic Delay, Prior Hour | | | | | | | | | 0.007 | *** |
| | | | | | | | | | (0.002) | |
| Traffic Complexity, Prior Hour | | | | | | | | | -0.028 | *** |
| | | | | | | | | | (0.001) | |
| Traffic Density, Prior Hour | | | | | | | | | 0.003 | |
| | | | | | | | | | (0.002) | |
| Coordination Load, Prior Hour | | | | | | | | | 0.001 | |
| | | | | | | | | | (0.001) | |
| Errors Made, Prior Hour | | | | | | | | | -0.020 | *** |
| | | | | | | | | | (0.004) | |
| *Control Variables* | | | | | | | | | | |
| Forecasting Work | 0.032 | *** | -0.020 | * | -0.006 | | -0.005 | | -0.003 | |
| | (0.009) | | (0.008) | | (0.008) | | (0.008) | | (0.008) | |
| Routing Work | -0.001 | * | -0.004 | *** | -0.005 | *** | -0.005 | *** | -0.005 | *** |
| | (0.000) | | (0.000) | | (0.000) | | (0.000) | | (0.000) | |
| ADS Use in Unit | 0.006 | *** | 0.018 | *** | 0.019 | *** | 0.019 | *** | 0.018 | *** |
| | (0.001) | | (0.001) | | (0.001) | | (0.001) | | (0.001) | |
| Fixed Unit Effects | included | | included | | included | | included | | included | |
| Fixed Day of Week Effects | included | | included | | included | | included | | included | |
| Fixed Hour of Day Effects | included | | included | | included | | included | | included | |
| $R^2$ | 0.110 | | 0.229 | | 0.302 | | 0.304 | | 0.322 | |
| N | 21930 | | 21930 | | 21930 | | 21930 | | 21930 | |

* $p < 0.05$
** $p < 0.01$
*** $p < 0.001$

## REFERENCES


Accenture. 2019. AI: Built to Scale. https://www.accenture.com/us-en/insights/artificial-intelligence/ai-investments. Accessed January 3, 2020.

Asch, S.E. 1956. Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs*, 70.

Bainbridge, L. 1982. Ironies of automation. *IFAC Proceedings Volumes*, 15(6), 129-135.

Bellet, T., Cunneen, M., Mullins, M., Murphy, F., Pütz, F., Spickermann, F., & Baumann, M. F. 2019. From semi to fully autonomous vehicles: New emerging risks and ethico-legal challenges for human-machine interactions. *Transportation Research Part F: Traffic Psychology and Behaviour*, 63, 153-164.

Cialdini, R.B. 2007. Descriptive Social Norms as Underappreciated Sources of Social Control. *Psychometrika* **72,** 263.

Endsley, M.R. 2017. From Here to Autonomy: Lessons Learned from Human-Automation Research, *Human Factors*, 59(1):5-27.

Fisch, J.E., Laboure, M., Turner, J.A. 2018. The Emergence of the Robo-Advisor, *Wharton Pension Research Council Working papers*, 10.

Gkartzonikas, C., & Gkritza, K. 2019. What have we learned? A review of stated preference and choice studies on autonomous vehicles. *Transportation Research Part C: Emerging Technologies*, 98, 323-337.

Kahneman, D. and D.T. Miller. 1986. Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93: 136-153.

JM Logg, JA Minson, DA Moore. 2019. Algorithm appreciation: People prefer algorithmic to human judgment, *Organizational Behavior and Human Decision Processes*, 151: 90-103.

McKinsey. 2019. Global AI Survey: AI proves its worth, but few scale impact. https://www.mckinsey.com/featured-insights/artificial-intelligence/global-ai-survey-ai-proves-its-worth-but-few-scale-impact, accessed January 7, 2020.

Rühr, A, Berger, B, Hess, T. 2019. Can I Control My Robo-Advisor? Trade-offs in Automation and User Control in (Digital) Investment Management. *Twenty-fifth Americas Conference on Information Systems*, Cancun, 2019.

Shiv, B., & Fedorikhin, A. 2014. Heart and Mind in Conflict: the Interplay of Affect and Cognition in Consumer Decision Making, *Journal of Consumer Research* 26 (3), 278–292.

Woods, D.D. & Cook, R.I. 2006. Incidents: Markers of resilience or brittleness? In E Hollnagel, DD Woods, & N Leveson (eds.) *Resilience Engineering: Concepts and Precepts* (pp. 69-76). Aldershot, UK: Ashgate.