EARLY INTERVENTION SYSTEMS

January 19, 2021

Zoe Russek and Dylan Fitzpatrick, University of Chicago Crime Lab

I. Executive Summary

Early intervention systems (EISs), sometimes called early warning systems, are proactive accountability tools which use police administrative data to try to identify police officers who are on a trajectory that may jeopardize either public safety or their performance on the job, then prescribe interventions to support the officer and reduce the risk of a future adverse event. An EIS provides a structured way for police departments to prioritize officer support resources by delivering an intervention to the officers most likely to benefit from support.

EISs have been widely used as instruments for police accountability since the U.S. commission on Civil Rights recommended that all police departments implement an EIS in 1981 (Walker, Alpert, & Kenney, 2001). In 1997, EISs also became a hallmark of police reform when the consent decree resulting from the U.S. Department of Justice's first pattern or practice investigation in Pittsburgh, PA included a requirement to implement an EIS (Davis, Henderson, & Ortiz, 2005). In jurisdictions operating under consent decrees, those agreements, and the independent monitors enforcing them, will often dictate the design of those systems, as well as what the system should consider "risky behavior" worthy of intervention. Many large US cities currently employ EISs, including Los Angeles, Chicago, Detroit, Seattle, and New Orleans.

The best available data are consistent with the potential for EISs to reduce adverse policing outcomes, although definitive evidence for the existence of those benefits (and their magnitudes) remains lacking:

- The logic of prioritizing supports and other interventions on specific police officers assumes that a small proportion of officers account for an outsized share of adverse policing outcomes. The best available evidence seems to support this hypothesis.
- EISs can only help prioritize supports if the *ex ante* risk of an adverse policing event is predictable in the administrative data that police departments routinely collect. This seems to be the case, although predictability of risk can vary across policing outcomes, over the life of an officer's career (greater predictability later in an officer's career), and by the nature of the statistical approach employed (formal predictive analytics do much better than simple 'rules of thumb').
- EISs can only prevent adverse policing outcomes if the supports that are being targeted are actually effective. There is encouraging evidence that effective supports do exist, although the degree to which existing EISs wind up employing the most effective supports rather than less-effective interventions remains unknown.
- Building on this last point, while EISs seem to have the potential to prevent adverse policing outcomes, there is little good empirical evidence on the actual impacts of the types of EISs that police departments have implemented on the ground in practice.

In what follows we first build out in greater detail the logic behind EISs and discuss the potential unintended consequences as well as hoped-for benefits. We then discuss some of the key design decisions that may affect the impacts that an EIS will have in practice, including what outcomes to focus on predicting, at what career stage, as well as the statistical methods that are employed to carry out that prediction exercise. We conclude with a review of the (quite limited) research literature of the effects of existing EIS systems on policing outcomes.

II. Conceptual Framework of an EIS

The primary goal of an EIS is to proactively support officers who, based on a history of previous incidents or job performance indicators, have demonstrated an elevated risk of future adverse events on the job. The outcomes in question can vary across departments but typically include a variety of adverse events that vary in severity, including:

- Excessive use of force
- Disciplinary action resulting in officer suspension
- Alcohol or other substance-related incidents
- Citizen-generated complaints of misconduct

Reducing adverse incidents occurring while *off-duty* is also often priority for departments, as officers who exhibit problematic behavior off the job, such as domestic violence or substance abuse, may directly benefit most from wellness supports and may be at heightened risk for misconduct on the job as well.

By targeting officers believed to be at elevated risk of an adverse event with some sort of an intervention, an EIS aims to get the officer on better footing and reduce the risk of future incidents, thereby lowering the frequency of misconduct for the department as a whole over time. Departments facing budget and human capital constraints cannot afford to intervene and provide additional supports (which can sometimes include fairly intensive supports) to all officers, so an EIS provides a structured approach for prioritizing supports where they may be most beneficial, as well as directing supervisory attention where it is needed most.

Because officers are targeted for intervention based on *predicted* risk of some future event based on administrative data, departments cannot simply terminate or otherwise punish officers flagged by an EIS. Instead, interventions are typically designed to be *supportive* rather than *punitive*. EISs exist outside of departmental disciplinary systems and are generally intended to improve an officer's wellbeing and preparedness for difficult situations. This is a common source of confusion with EIS systems, which can easily get conflated with efforts to improve officer disciplinary systems that are designed for a very different task – holding officers accountable after an adverse event has occurred.

The success of an EIS in reducing officer risk of misconduct relies on the effectiveness of the interventions being targeted. Interventions targeted by an EIS differ across departments and even across officers within a department:

- An intervention may be as simple as a conversation with a supervising officer, in which the supervisor addresses previous incidents and discusses a plan for handling them better in the future.
- Alternatively, some officers may be identified as needing additional training, such as for proper use of force or for better situational awareness in high-stress settings.
- Others may be identified as having untreated mental health challenges and connected with therapists trained to work with officers to improve mental health and wellness.

The appropriate support likely depends on the specific set of challenges faced by an officer, as indicated by the officer's history within the department and relationship with his or her supervisor.

While an EIS is intended to reduce the frequency of unfavorable events and improve job performance among officers, the possibility exists for these systems to be ineffective or to have negative effects due to failures of design or implementation. For example, a police department may institute an EIS that identifies high-risk officers but doesn’t actually deliver supports to those individuals, resulting in no change to policing outcomes. Even if supports are delivered to officers, they may be ineffective or have negative consequences. (Most social programs for the civilian population are found not to work, so there is the logical possibility the same could be true for supportive programs for police officers as well). Some observers have noted that implementation of an EIS could as a logical matter also potentially have the unintended effect of reducing officer engagement with policing activities to avoid being flagged for intervention (Worden, 2013). Finally, there is the possibility that being flagged by an EIS could induce feelings of persecution in officers already on uneven footing, thereby increasing job-related stress and associated risk of future adverse events. This was the purported reason for the Chicago Police Department discontinuing one early EIS tool after only two years, as officers felt the system unfairly targeted officers and led to adversarial questioning (USDOJ, 2017). The ultimate effect of any given EIS system in practice is an empirical question, one that is also likely to depend on the exact design of that specific system.

III. Identifying Officers for Intervention

The ability of departments to effectively direct officer supports relies on first having a system for identifying high risk individuals. The most common approach is to flag officers based on a set of thresholds across key performance indicators. An EIS traditionally uses three types of rule-based thresholding to identify officers (Walker, Milligan, & Berke, 2006).

- *Department-level* thresholds are those set at the department level and affect all officers, regardless of position (e.g., an officer has two or more citizen complaints or three or more use-of-force reports in the previous 90 days).
- *Performance indicator ratios* are thresholds set by ratios between different performance indicators (ex: ratio of use-of-force reports to arrests).
- P*eer-officer average* thresholds compare officers to what is typical among others with similar assignments and positions.

A rule-based threshold system is relatively low-cost to produce, as thresholds can be determined from straightforward examination of performance indicators that are already monitored by police supervisors, but they may fail to identify subtle indicators of risk that are present in administrative records. An improvement over these systems is to use statistical analysis of the entire set of available data variables to develop statistical models that are specifically designed to predict future adverse events.

In a case study based on administrative records from a metropolitan police department in the U.S., high-risk officers identified by prediction models went on to experience adverse events (including suspensions, sustained complaints, or off-duty domestic or substance use complaints) at a rate 1.7 times higher than those identified by a rule-based EIS. Table 1 shows adverse event rates in the top 1% of officers identified by a rule-based system, the top 1% identified by prediction models, and in the population of all officers. The top 1% of officers identified by prediction models had an adverse event rate of 50 events per 100 officers within two years, or one event for every two officers flagged by the EIS. Prediction models are able to do a better job of identifying high-risk officers when compared to the rule-based system, by finding statistically relevant relationships across time periods and specific categories of misconduct.

Accurately identifying officers with higher-than-average adverse event rates is critical for maximizing benefits from interventions targeted by an EIS. Consider a simple example in which 10 officers out of 100 will go on have an adverse event. If a department were to flag 10% of all officers at random, only one out of those ten flagged officers would have gone on to have an adverse event on average. Even with a perfectly protective intervention that reduced risk of that event to zero, the randomly targeted system would only prevent a single event. In contrast, with a perfectly predictive model, all ten of those flagged would have had the event, and the system would prevent ten times as many adverse events for the same number of officers delivered an intervention. This helps illustrate why the share of people at the top of the predicted risk distribution who would go on to have had that outcome is such a common measure of how well a statistical prediction model is working, and why making that share as large as possible is such an important objective from the perspective of averting adverse events.

Table 1. Comparison of adverse event rates in the top 1% of officers flagged by rule-based vs. prediction-based EISs in a large U.S. police department.

	Rate of Adverse Events (Per 100 Officers)
All Officers	11
Officers Identified by Rule-Based EIS	29
Officers Identified by Prediction Models	50

An EIS based on predictive analytics has two advantages over a rule-based EIS: first, performance indicators are selected through analysis of historical data to determine which specific factors are most predictive of adverse outcomes on the job; second, information is more readily integrated across multiple risk factors to identify individuals in need of intervention. As a result, an officer who is slightly below all thresholds across multiple performance indicators would be missed by a traditional rule-based EIS, while an EIS based on statistical analysis may correctly identify that officer as high risk.

Helsby et al. (2017) report predictive performance of models designed for an EIS developed for the Charlotte-Mecklenburg Police Department. The study compares accuracy of a new, model-based EIS to an existing, rule-based EIS already in place at the department and finds that the new system provides a 20% reduction in false positives and a 75% increase in true positives among officers identified by the system. As these results indicate, statistical prediction models provide a promising new approach for identifying officers most likely to benefit from support designed to mitigate risk, but these tools have yet to be widely adopted in operational EISs and additional studies are needed to fully evaluate their performance.

The performance of such statistical models for risk prediction also is likely to rely heavily on several key design choices. For the most serious types of officer transgression (e.g., improper use of force), every individual case is extremely troubling from a public policy and civil rights perspective, but from a narrow statistical perspective there can sometimes be too few such cases in a statistical sense to enable the construction of an accurate predictive model. Targeting more frequent proxy measures for misconduct, predicting some aggregate of multiple outcomes, or combining data across multiple departments may all improve predictability, especially for departments with fewer officers and less historical data for training models.

Past work with such data also suggests that prediction performance can often be improved by using detailed information on the time frames in which previous events occur; for example, several complaints occurring within a single month may indicate that an officer is going through a particularly difficult period, while a pattern of complaints over multiple years could be evidence of a more ingrained issue with the officer's training. Previous statistical analysis also suggests that past adverse events that happened relatively more recently are more predictive of future risk than those that happened relatively longer in the past. So for the design of a statistical model, constructing candidate predictors that are sensitive to when past events happened may be more predictively useful than models that just ask whether an officer had such a past event ever.

Predictability of risk for a given officer seems to improve over the course of their career as a department collects additional complaints or use of force records, which provide important signals that an officer may be facing difficulties on the job. Statistical models detect meaningful interactions between these signals, connecting patterns over multiple time horizons and specific categories of misconduct to identify patterns that are predictive of problem behavior. Alternative approaches to identification of risk may be needed for early career officers with shorter histories of departmental data records, as they may have too little data for models to accurately assess risk.

Table 2 shows base rates of several adverse event categories among sworn officers in one large metropolitan police department in the U.S., as well as some key performance metrics from prediction models[1] trained to identify the top 1% of officers most likely to experience these events. Dividing the share of officers identified by the model who actually went on to experience the adverse outcome (known as *precision*) by the outcome *base rate* (the prevalence of the outcome in the department as a whole), the resulting ratio (called the *predictive lift*) provides a useful indicator of a model's predictive ability. For prevention the goal is to predict and flag a set of officers among whom prevalence of the adverse outcome would have been as high as possible, so that supports and interventions can prevent as many of these outcomes as possible.

Table 2 highlights some common features of these types of predictive models. First, it is typically easiest to achieve a high rate of precision (high share of officers who are flagged would have gone on to have some outcome) for the most common outcomes, such as complaints. For example, for the officers with highest 1% of predicted risk of future citizen complaints, officers went on to actually receive a complaint at a rate of 69 per 100 officers, a rate 2.3 times higher than a system that randomly targets officers for intervention (which would be the base rate).

More serious adverse event categories, such as complaints of excessive force with unnecessary physical contact, complaints resulting in suspension, and complaints of domestic violence, can be more difficult to predict due to these events being relatively infrequent. Still, models provide predictive lift of 7.7, 3.2, and 9.1 respectively on these higher priority categories.

Table 2. Performance of statistical prediction models on predicting adverse event categories using officers' most recent 5 years of administrative records, where officers with top 1% of predicted risk are targeted for intervention.

	Base Rate	Precision	Lift	Recall
Any citizen complaint	0.301	0.69	2.3	0.02
Excessive force complaint with physical contact	0.005	0.02	7.7	0.08
Complaint with suspension	0.005	0.02	3.2	0.03
Domestic violence complaint	0.002	0.02	9.1	0.09

Statistical models based on administrative data can provide a predicted risk score for each individual officer, but departments then face a key policy decision: about how many of the highest-risk officers should be flagged by an EIS? This decision presents departments with a tradeoff. All else equal, the smaller the share of officers who are flagged, the larger the share of those who are flagged would have gone on to have the adverse event (and hence the more adverse events prevented per dollar of intervention spending, all else equal). On the other hand,

[1] For this case study, we trained gradient-boosted decision trees to predict whether an adverse event occurs in the two years following the point of prediction, using up to five years of officer-level citizen complaints, use-of-force records, and departmental disciplinary actions as inputs to the prediction models.

the smaller the share of officers who are flagged, the smaller the share of total adverse events that would happen in the department will be prevented by the EIS. More support resources are needed as more officers are flagged by the system, but the benefits from support may be diminished for officers with lower predicted risk.

To get some sense of the issue, in Table 2 we also report the share of total events that would happen within this department for each adverse outcome that would have occurred among the top-predicted 1% (a measure called *recall*). If the department were to flag 1% of officers for intervention, prediction models would correctly identify 2-9% of all officers who went on to experience adverse events across event categories. Increasing the share of officers who are flagged increases the share of total events in the department that would have occurred among the flagged officers, but the total share increases as the share flagged increases at a decreasing rate. Ultimately the decision about where the risk threshold should be for flagging officers, and hence how many officers will get flagged, is a policy question not a statistical question.

IV. How to Intervene?

An EIS is only as effective as the specific support that it is intended to deliver. Proof-of-concept evidence suggests that effective supports are possible, although the degree to which actual EIS systems implemented in the real world rely on effective vs. less effective supports is unclear.

Supervising officers represent the first line of defense for mitigating risk in an officer that has been flagged by an EIS. Owens et al. (2018) show that the right kind of structured supervisor-officer conversations have the potential to be helpful. Owens and her team conducted a randomized trial at the Seattle police department that assessed effects of a supervisory program in which officers and their supervisors discussed a recent encounter that the officer had with a public citizen. Officers were prompted to consider their thought processes and actions, while supervisors applied principles of procedural justice in talking through the encounter. That is, there was some clear intentional structure to these conversations, informed by the field's understanding of behavioral science and behavior change. The study found that following these conversations, officers were 12% less likely to resolve incidents with an arrest and were between 16 to 50% less likely to be involved in use of force incidents, with officers working in highly active beats demonstrating the largest reductions in arrests. Structured conversations with a supervisor represent a low-cost intervention that has the potential for significant reductions in adverse outcomes, particularly when accurately targeted toward officers with a high propensity for complaints or use of force.

Studies on officer wellness indicate high rates of untreated mental health problems among police officers. A survey of officers in the Dallas police department found that 26% of officers reported current symptoms of mental illness, only 17% of whom had pursued mental health care in the previous year (Jetelina, 2020). The events of 2020 may have increased job-related stress among police officers. As noted in a separate CCJ report on Officer Wellness, few empirical studies have examined the causal effects of mental health services on officer and public safety outcomes. Still, provision of services such as therapy provided by clinicians embedded at a police department or through telehealth services represent a promising support for officers experiencing job-related stress and anxiety. A 2018 review of published literature on the use of telemedicine to provide mental health services found consensus that these services are as effective as in-person interventions for both diagnosis and treatment across a range of mental health conditions (Shigekawa et al., 2018). Still, evaluation of these interventions often relies on academics examining the best, experimentally validated telehealth services from among tens of thousands of providers and applications available to anyone with a smartphone. Departments need to be discerning in what sorts of mental health supports they provide to ensure that the intervention is effective in practice.

Lastly, training sessions targeted towards high-risk individuals provide another approach to intervening on multiple officers simultaneously, but the effectiveness of different types of training and training modalities remains uncertain. We discuss in detail the empirical literature on effects of different forms of training in separate CCJ reports on De-Escalation and Use of Force Training, Procedural Justice Training, Implicit Bias Training, and Training Modalities.

V. The Landscape of EIS across U.S. Police Departments

Our discussion so far suggests that the impact of an EIS will depend on policy design decisions like what outcome to predict, how that outcome is predicted (formal statistical model vs. rule-based thresholding), what share of officers to flag, and what supports to provide. This means that while EIS systems may have the potential to prevent adverse outcomes, the actual impact of EISs as implemented in the real world will depend on the specific design decisions made.

The possibility for significant variability in design is apparent when looking at EISs that have been implemented to date in departments across the U.S. Most modern EISs utilize a rule-based thresholding approach to identify officers who would most benefit from some form of intervention. Departments first select a set of individual performance indicators (e.g., uses of force, citizen complaints, and departmental policy violations), then consider the number of these incidents a given officer has in comparison to other department members. For each performance indicator, counts above some threshold value generate an automatic alert in a computer application notifying the department that the officer needs an intervention of some kind. While there is some variation in the performance indicators used in various jurisdictions, nearly all EISs include a count of officer-involved uses of force, officer-involved shootings, and citizen-initiated complaints. See Table 3 for a summary of the 10 most common performance indicators in EIS across eight U.S. cities.

Most rule-based systems will involve multiple types of thresholds for evaluating officers, such as those outlined in Section III above. Pittsburgh was the first city to implement a peer grouping system to identify officers and this is now common practice around the U.S., including in the Los Angeles Police Department’s TEAMS II system implemented as part of the 2001 consent decree. In peer-officer average systems, officers are compared to a peer-group-specific threshold that is set to the value two standard deviations from the average among their own peer group. For example, if an officer receives their fourth complaint, the computer application compares that number to the average number of complaints within their peer group, alerting the supervisor if four is more than two standard deviations above the peer group average.

Table 3. Common Performance Indicators from EIS in eight U.S. cities.

10 Most Common Performance Indicators	**Detroit**	**New Orleans**	**Pittsburgh**	**Los Angeles**	**Cleveland**	**Ferguson**	**San Diego**	**Seattle**
Officer- involved use-of-force incidents	x	X	x	x	x	x	x	x
Officer-involved shooting incidents	x	X	x	x	x		x	
Canine bites	x	X	x	x	x			
Citizen-initiated complaints	x		x	x	x		x	x
Citizen or department-initiated commendations or awards	x	X	x	x	x	x		
Departmental disciplinary actions	x	X	x	x	x	x		
Training and reassignment history	x	X	x	x		x		
Officer-involved civil suits or administrative claims	x	X	x	x	x	x		
Possible or reported officer misconduct	x	X	x	x	x	x		x
Vehicle pursuits and car crashes	x	X			x		x	

More recently, organizations such as the University of Chicago Crime Lab, Data Science for Social Good, and Benchmark Analytics have developed EISs that employ statistical methods to identify specific indicators that are predictive of future adverse events on the job. As discussed above in Section III, statistical models can be trained to identify the small set of officers at highest risk of future adverse events, relying on the entire history of administrative records that a department has collected on each officer.

The existing research literature is extremely limited on direct, causal effects of EISs on public safety and officer performance outcomes in practice. The studies that do exist primarily rely on before-after comparisons, including a study by Walker, Alpert, & Kenney (2001) which played a key role in the proliferation of EISs around the US. The study showed a decrease in use of force and citizen complaints for officers identified by an EIS one year after identification in three U.S. jurisdictions. Results from subsequent before-after studies are mixed, with some studies similarly observing reductions in citizen complaints (Macintyre et al., 2008; Briody and Prentzler, 2020), and others finding no statistically significant effects following EIS implementation (Wordern, et al. 2013; Shjarback, 2015). Problems with research design unfortunately restrict the ability of these studies to make causal claims with any certainty.[2] Most notably, comparing outcomes before and after implementation prevents these studies from disentangling the effects of the EIS with the effects of broader reforms that may be occurring in these departments at the same time. Observed reductions in use of force, citizen complaints, or other adverse outcomes that coincide with the adoption of an EIS could be a result of other changes in supervision, accountability, or training. (That is, it is possible that whatever internal or external reasons caused a department to adopt an EIS at that time may have led to other changes within the department that also serve as candidate explanations for any observed change in policing outcomes).

The limited EIS literature indicates that further, more robust exploration is required to fully assess whether EISs bring benefits to policing outcomes as intended. Principally, issues with experimental research designs restrict the ability of existing studies to make causal claims with any certainty. Standardization of reporting and evaluation is needed to better compare EISs across cities, particularly with respect to methods used to identify high-risk officers and the interventions that are ultimately targeted by an EIS.

In this report, we have outlined several necessary building blocks of a successful EIS and highlighted suggestive evidence that these components can be effective in isolation. While direct effects of EISs on policing outcomes have not been evaluated rigorously, evidence shows that prediction models can successfully predict *ex ante* risk of adverse policing events to some extent, and officer supports can reduce this risk when implemented properly. Additional research is needed to assess the integration of these components, and to inform future development of EISs as tools for accountability and support.

References

Briody, M., & Prenzler, T. (2020). The New Zealand police early intervention system: A review of implementation and impact issues. *International Journal of Police Science & Management*, 22(3), 297–307.

Davis, R. C., Henderson, N. J., & Ortiz, C. W. (2005). *Can federal intervention bring lasting improvement in local policing? The Pittsburgh consent decree*. New York, NY: Vera Institute of Justice.

Helsby, J., Carton, S., Joseph, K., Mahmud, A., Park, Y., Navarrete, A., Ackermann, K., Walsh, J., Haynes, L., Cody, C., Patterson, M. E., & Ghani, R. (2018). Early intervention systems: Predicting adverse interactions between police and the public. *Criminal Justice Policy Review*, 29(2), 190–209.

Jetelina K.K, Molsberry R.J., Gonzalez J.R., Beauchamp A.M., & Hall T. (2020). Prevalence of mental illness and mental health care use among police officers. *JAMA Network Open*, 3(10).

[2] See the Methods for Research Review report for a summary of the criteria used to assess the methodological rigor of existing research and to determine what studies constitute reliable evidence in the present report.

Macintyre, S., Prenzler, T., & Chapman, J. (2008). Early intervention to reduce complaints: An Australian Victoria police initiative. *International Journal of Police Science & Management*, 10(2), 238-250.

Owens, E., Weisburd, D., Amendola, K., & Alpert, G. (2018). Can you build a better cop?: Experimental evidence on supervision, training, and policing in the community. *Criminology & Public Policy*, 17(1).

Shigekawa, E., Fix., M., Corbett, G., Roby, D. H., & Coffman, J. (2018). The Current state of telehealth evidence: A rapid review. *Health Affairs*, 37(12).

Shjarback, J. A. (2015). Emerging early intervention systems: An agency-specific pre-post comparison of formal citizen complaints of use of force. *Policing: A Journal of Policy and Practice*, 9(4), 314-325.

USDOJ. (2017). *Investigation of the Chicago Police Department.* US Department of Justice, Civil Rights Division.

Walker, S., Milligan, S. O., & Berke, A. (2006). *Strategies for intervening with officers through early intervention systems: A guide for front-line supervisors*. US Department of Justice, Office of Community Oriented Policing Services.

Walker, S., Alpert, G. P., & Kenney, D. J. (2001). *Early warning systems: Responding to the problem police officer*. Washington, DC: US Department of Justice, Office of Justice Programs, National Institute of Justice.

Worden, R. E., Kim, M., Harris, C. J., Pratte, M. A., Dorn, S. E., & Hyland, S. S. (2013). Intervention with problem officers: An outcome evaluation of an EIS intervention. *Criminal justice and behavior*, 40(4), 409-437.