URBANLABS
CRIME LAB

Harnessing Data to
Advance Justice and
Reduce Gun Violence

crimelab@uchicago.edu
773 834 4292

33 North LaSalle Street
Suite 1600, Chicago, IL 60602

urbanlabs.uchicago.edu/labs/crime

Workforce Allocation Summary

The allocation of police department resources in most US cities is based on the desires and intuition of key decision-makers and often winds up being highly political and unequal. This opaque and ad-hoc process can also be slow to reallocate resources in response to changes in community needs, localized crime spikes, or changing incident patterns. Too often, officers aren't available when and where they're most urgently needed. This inequitable distribution has real consequences, as we can see in our own home city of Chicago. In some areas of the city, residents receive rapid responses to both emergency and non-emergency 911 calls; in other areas, there are no officers available to respond for hours to 911 calls, sometimes even for violent incidents like robberies or shootings.[1] These discrepancies also undermine officer morale as some officers are consistently overworked while others are underutilized.

To improve efficiency, equity, and transparency in patrol staffing, the University of Chicago Crime Lab developed a data-driven approach to re-deploy officers to the busiest parts of the city. By allocating additional officers to areas of the city that are struggling to adequately respond to calls for service, we aim to reduce the overall imbalance in 911 responses across Chicago neighborhoods.

The remainder of this memo describes the details of our approach.

Methodology

The Crime Lab, in response to a request from CPD leadership, developed a planning model to help the department better allocate officer resources. While the Crime Lab's model was developed initially with the Chicago context in mind, in principle the approach can be used anywhere.

The Crime Lab model was designed to optimize deployment by focusing on a key metric: the amount of time officers spend responding to 911 calls compared to the amount of time officers spend doing other work. CPD's stated goal was to have this equalized across areas of the city. By spreading officer workload more evenly, both the speed with which officers are able to respond to calls and the amount of time they are able to spend addressing each call may improve.

Every police department has targets for the percentage of time officers should spend answering calls and the percentage of time they should have left over to spend on administrative tasks and proactive policing. This percentage allocation is, ultimately, a policy decision, not a research question. CPD set its preferred ratio at 60/40. The Crime Lab developed a model to optimize patrol staffing across districts so that officers ideally spend no more than 60% of their time answering calls for service, freeing up 40% of their time to devote to other police work (administrative work,

[1] https://chicago.suntimes.com/2021/7/16/22558978/nonemergency-calls-alternate-response-section-chicago-police-department-kiama-doyle; https://cwbchicago.com/2021/08/two-hour-response-time-for-robbery-no-cops-available-for-a-man-shot-on-the-street-detectives-pulled-from-cases-to-patrol-lollapalooza-welcome-to-chicago-2021.html

UCHICAGO URBANLABS CRIME LAB

Harnessing Data to Advance Justice and Reduce Gun Violence

crimelab@uchicago.edu
773 834 4292

33 North LaSalle Street
Suite 1600, Chicago, IL 60602

urbanlabs.uchicago.edu/labs/crime

community policing, and proactive policing). Our unit of analysis is patrol cars, which is the equivalent of one or two officers, depending on the time of day.[2]

The Crime Lab model uses an iterative process to allocate patrol cars to support busy officers until there are no more overworked officers and every officer is able to spend no more than 60% of their time answering calls. In brief, the Crime Lab algorithm:

- Uses historical data to calculate how much time each individual patrol car assigned to each beat and shift spent responding to calls
- Sorts the patrol cars from busiest (i.e., those with the highest percentage of time spent responding to calls) to least busy
- Assigns an additional patrol car to reduce the workload of the busiest patrol car, assuming this will cut the workload of the busiest car in half
- Repeats the process until all officers are below the specified threshold of 60/40

The fastest way to equalize 911 response times across the city would be to reallocate existing staff as needed across the city. In practice, collective bargaining agreements and other institutional issues may constrain the ability to reallocate existing officers across areas or shifts. In that case, the model could also be employed to allocate new CPD staff as they join the workforce, which would then require more time to equalize 911 response times across areas.

Research Complications

Conceptually, the workforce allocation model is straightforward. However, implementing the conceptual framework behind the model using real-world data raises a number of complications:

- *Measuring time spent responding to calls isn't straightforward, so we use median on-call time as our measure of on-call time*. Given the possibility for data errors and inconsistencies inherent to any complex administrative dataset, we substitute observed on-call time with the median on-call-time by district, shift, and priority level in order to reduce the influence of any errors and outliers in the data.[3]

---

[2] Depending upon the police shift in question, a car may contain one or two officers. Second shift cars typically have one officer while 1st and 3rd shift cars typically have two officers. Most cars (excluding a few "floater" tactical units) are assigned to a specific beat within a district, but the model focuses specifically on cars rather than beats because certain beats are busy no matter how many officers are staffed there. A beat that receives 100 calls a week with one car will still receive 100 calls a week with two cars, but those calls will be more evenly spread out over two cars than one. Likewise, the model doesn't focus on individual officers because officers can switch between different cars and beats day-to-day and are more difficult to track. Concentrating on car call time allows for much more precise resource allocation.

[3] For example, imagine that Car A responded to 5 calls that were in District 1, Shift 1, and call priority level A1 and that those calls had recorded on-call times of 20 minutes, 40 minutes, 1 hour, 1.5 hours, and 6 hours. The raw total on-call time for that district, shift, and call priority combination would be 8.5 hours (and likely reflect data-entry error given the large outlier). Under our methodology, each on-call time falling under that particular

URBANLABS CRIME LAB
UCHICAGO

Harnessing Data to Advance Justice and Reduce Gun Violence

crimelab@uchicago.edu
773 834 4292

33 North LaSalle Street
Suite 1600, Chicago, IL 60602

urbanlabs.uchicago.edu/labs/crime

- *We assume that adding an extra car will cut the work of the original car in half, but that may not always happen in practice.* This assumption implies that if the model allocates an additional car to a beat, there is perfect coordination between both cars; for instance, if there are 10 calls per shift, each car would handle 5 calls. However, it's conceivable that both cars would respond to the same high priority calls rather than "dividing and conquering." It's also possible that even if cars perfectly split the number of calls, idiosyncratic variation in the type and length of call handled by each car would still lead to disparities in total call time between cars.
- *We assign additional patrol cars directly to the busiest car (within each beat and shift) rather than to a district overall.* While deploying an additional car at the district level might have some proportional impact on reducing the load on the busiest cars, it is unlikely to have the same concentrated impact on workload as directly assisting an overwhelmed car that is assigned to a given beat and shift. Ultimately, the deployment of additional cars is at CPD's discretion and may result in cars being assigned as more general district resources.
- *We assume that adding an extra car does not change the volume of 911 calls, but there are scenarios in which extra cars could either increase or decrease calls.* If residents frustrated by slow response times are suppressing the number of calls they would make otherwise, then adding cars and therefore improving response times could increase the volume of 911 calls overall. On the other hand, if additional police presence has some deterrent effect on crime, then an increase in the number of cars could reduce call volume.

Workforce allocation is a complex exercise, and because the process has not been data-driven for so long, there are many factors at play about which little is known. Some of these open questions may not be answerable without actually implementing the allocation model, which wound up not happening in the city of Chicago. Our model and policy simulation represent first passes at optimal staffing allocation. Actual implementation would ideally involve deployment of the base model's recommendations, accompanied by measurement, evaluation, and implementation of any refinements necessary to accomplish the stated goals of the city in question.

For more information, please contact Roseanna Ander at rander@uchicago.edu.

---

combination of district, shift, and priority level (1/1/A1) would be replaced by the median of 1 hour, leading to a total adjusted on-call time of 5 hours—1 hour per call. The reason the median is based on calls in each combination of district, shift, and priority level (rather than a global median of all calls) is that some types of calls in particular districts or times of day take longer than others; high priority calls may also have shorter response times if officers are able to get there faster, but they may also take longer to resolve. All these nuances are best captured by the median within each district, shift, and priority level combination.