ONLINE APPENDIX: Professional Motivations in the Public Sector: Evidence from Police Officers

Aaron Chalfin*   Felipe Goncalves†

*University of Pennsylvania and NBER, achalfin@sas.upenn.edu
†University of California, Los Angeles and NBER; fgoncalves@ucla.edu.

Contents

A Data Construction

Section II describes the various data sources we use and how we link them and construct our analysis sample. We provide more information here on the sample construction and specifically how many observations are dropped when restricting our raw data to our final analysis sample.

There are 5469 officers who work in the Dallas Police Department during our sample period. 685 (12.5%) of these officers are dropped because they never appear in the overtime data. An additional 2354 (43.0%) officers are dropped because they never work a patrol division assignment, never work two adjacent patrol division shifts, or because they are never observed taking a 911 call during their regular shift. These restrictions leave us with 2430 officers in our analysis sample.

Our raw 911 data consists of 2144492 calls made during our sample period. 171924 (8.0%) of these calls are dropped because the listed responding officers never appear in the overtime data working a patrol shift. 680479 (31.7%) of calls have a responding officer who does work patrol shifts but the call is outside the hours or dates we know the officer is working a regular shift. Note that these drops can occur for several reasons. These calls may be occurring during an additional shift taken by an officer, in which case we do not know the hours of the shift. They may also be taken during a regular shift for that officer but occur between adjacent overtime payments that list different regular shifts, and we are unable to infer the shift times for that day. These sample restrictions leave us with 1394380 calls in our analysis sample.

Our raw arrest data consists of 149339 arrests made during our sample period. 39034 (26.1%) of these arrests are dropped because the listed officers never appear in the overtime data working a patrol shift. 69083 (46.3%) of arrests have an arresting officer who does work patrol shifts but the arrest is outside the hours or dates we know the officer is working a regular shift. Similar to our dropped calls, these arrests may be occurring during an additional shift taken by an officer, in which case we do not know the hours of the shift. They may also be taken during a regular shift for that officer but occur between adjacent overtime payments that list different regular shifts, and we are unable to infer the shift times for that day. These restrictions leave us with 68569 arrests in our analysis sample.

Our sample restrictions focus attention on the set of officers who are primairily responsible for 911 call responses. The officers in our sample respond to an average of 1706.7 calls during our sample window. The officers who are not in our sample because they appear in the overtime but not working an 8-hour patrol shift respond to an average of 207.2 calls. Among the 685 officers who never appear in the overtime data, only 29 ever appear responding to a 911 call, and this subset respond to an average of 27.7 calls.

Sample Construction Robustness

We present here several analyses to validate that our main findings are robust to different choices in our sample construction.

Our sample construction approach requires dropping dates between adjacent overtime payments where officers have differing shift schedules. We assess whether there might be sample selection issues induced by disagreeing shift schedules by restricting the sample to officers who only have one regular shift during the entire sample window.

This analysis is reported in Figure A-9. All of the essential features of our principal findings remain present. Arrests fall throughout these officers' shifts and the rate of court conviction increases. These patterns in this restricted sample provide some assurance that the officers who are excluded due to uncertainty about their work shifts are not appreciably different and their exclusion does not alter the findings.

We next examine the role of the dropping of a sizable share of calls and arrests which do not occur on officers' regular shift times.

First, we ask how our analyses change when we restrict attention to officers where at least 75% of their calls match to their regular shift, presented in Figure A-10. The main data patterns, a within-shift decline in arrest propensity and increase in court convictions, persist in this subsample of the data.

Second, we ask how our analysis would change if we *expanded* our sample to irregular shifts. Our current data construction approach is designed to only use the available information on regular shifts, since there is no official information available on irregular shifts taken by officers. However, we can use a data-driven approach to infer officers' shifts on these days.

We do so as follows: we match officers' regular shift days to the 911 data, and we only keep calls that do not match (i.e. are currently dropped). For each officer, we calculate the eight most common hours of the day when they are taking calls, and we let the first of these hours be the start of their "irregular shift." We then construct 8-hour shifts for those days when officers are taking 911 calls outside their regular shifts. This approach captures 63% of all calls taken on non-regular shift days. While this approach is ad-hoc and thus not our preferred practice for the main analysis sample, these data-inferred irregular shifts can provide information on whether the lack of information on these shifts induces sample selection bias.

Figure A-11 presents our main analyses for this sample of irregular shifts. We find an identical pattern of arrests, with a low first-hour arrest rate that increases in the second hour, followed by a decline throughout the shift. The regression only controls for officer fixed effects, since we do not observe an officer's assigned division on irregular shift days. For court convictions, we continue to find no decline in convictions near the end of the shift. The pattern appears flat rather than increasing as in our baseline sample, though our standard errors are larger. We conclude from this analysis that our findings are not biased by restricting attention to regular shifts.

Next, our sample restricts to shifts when officers take at least one 911 call as an additional confirmation that our shift construction process captures days when officers are working patrol. Figure A-12 removes this requirement, and the arrest and court conviction patterns are identical to our baseline findings.

One additional concern we confront is the possibility that our sample construction leads to mismatched calls and arrests that are occurring on the same day as an officer's regular shift but on hours just outside their official patrol hours. For example, if officers take several hours to process an arrest and then place the final hour as the arrest hour rather than the original arrest time. To probe this concern, we expand our sample to include the four hours before and four hours after every regular shift, and we re-plot the raw arrest and court conviction time paths, presented in Figure A-13. Reassuringly, a very small share of arrests are occurring before or after the regular shifts, providing a validation that our regular shift time stamps are capturing when officers are patrolling. In addition, the patterns of declining arrests and increasing court convictions persist when examining this expanded shift window.

B Do Some Officers Engage in "Collars for Dollars"?

While our principal findings run contrary to the narrative that officers make additional low-quality arrests at the end of the shift in order to receive overtime pay, it is still possible that a fraction of officers engage in this practice even if the behavior of these officers is not detectable in the aggregate data. In this section, we investigate heterogeneity across officers in their late-shift arrest behavior, focusing on the officers who are especially likely to concentrate their arrest activity at the end of their shift. We assess whether the late-shift arrests made by these "late-arresters" are of lower quality, are more likely to be officer-initiated or disproportionately target minority citizens.

We begin by investigating whether officers, in fact, differ systematically in their propensity to make late-shift arrests. Given that there is evidence that police officers do differ systematically in their overall propensity to make arrests (Weisburst, 2022), any analysis that evaluates differences in end-of-shift arrests must account for overall differences in arrest propensity. Accordingly, we run a regression in which the unit of analysis is a given arrest and calculate whether the arrest i occurs in the last two hours of the arresting officer's shift:

$$Late_i = \alpha_{o(i)} + \phi_{dwh} + \theta_{dws} + \xi_{dm} + \epsilon_i \tag{A-1}$$

In (4), we are evaluating officer differences in the likelihood of making a late arrest conditional on overall arrest activity. As with our baseline analysis in Equations (1) and (2), our fixed effects include division $\times$ day-of-week $\times$ hour, division $\times$ day-of-week $\times$ shift, and division $\times$ year-month. The objects of interest are the set of officer-level fixed effects, $\alpha_{o(i)}$, which document systematic differences in the share of arrests that are made at the end of the day. Notably, because each officer makes a finite number of arrests, each fixed effect will be estimated with error and naturally some fixed effects will be estimated greater precision than others. To adjust the distribution for estimation error, we use a Bayes shrinkage approach similar to that employed in the teacher value added literature (Morris, 1983).

The results of this regression are presented in Figure A-14 which plots both the unadjusted and shrunken distributions of the officer fixed effects.[1]

We next use the shrunken fixed effects to explore whether the officer fixed effects are correlated with several signature behaviors of the collars for dollars story. In particular, we ask whether, among officers who are "late arresters," early versus late-shift arrests differ with respect to 1) the probability of conviction, 2) the share of arrests that are officer-initiated, 3) the share of arrests that are of African-American suspects and 4) the types of arrests that are made. To the extent that late arresters are differentially likely to make low quality arrests, officer initiated arrests, arrests of African-Americans or arrests that are more likely to lead to overtime pay, this potentially forms the basis for a claim that such behavior is motivated by the desire to secure access to overtime pay.

[1]Because of the presence of other fixed effects in the regression, the officer fixed effects are approximately centered at zero. The fixed effects are not exactly centered at zero because officers have different numbers of arrests, and the fixed effects distribution is at the level of the officer. The average of the fixed effects weighted by the number of arrests is equal exactly to 0.

We explore these relationships in Figure A-15 which, for each outcome, plots the mean of the dependent variable separately for early shift arrests (the dashed line) and late-shift arrests (the solid line) for officers above a given percentile of the distribution of the shrunken fixed effects. We fail to see evidence that, among the late arresters, early and late-shift arrests differ with respect to the share of arrests they make that are officer-initiated and the share of arrestees who are non-white. With respect to conviction rates, if anything arrests during the final four hours of the shift have a higher conviction rate. We also generate a predicted overtime variable by multiplying each arrest by its expected number of overtime hours using statistics presented in Table A-1. For each officer we compute predicted overtime hours on the basis of the distribution of that officer's arrest charges. There is little evidence that late arresters are differentially likely to make the types of arrests that are more likely to lead to overtime hours. in their shift. Had this been the case, we would have expected the dashed and solid lines to have switched positions at the top of the distribution.

Taken as a whole, the evidence suggests that while some officers consistently make late-shift arrests, this behavior cannot be explained by the type of strategic behaviors that have been suggested as pillars of the collars for dollars story. Instead, it is possible that officers who tend to make late-shift arrests simply have greater ability or skill to make late-shift arrests, either because their skills erode over the course of a shift at a lower rate or for some other reason.

C Robustness Analyses

We argue that the observed patterns of reduced arrest frequency and increased severity of arrests made near shift-end reflect officer preferences and, in particular, officers' aversion to overtime work. In this section we test the validity of our empirical design and present a series of tests for whether our results are driven by officer preferences rather than an artifact of either the data or institutional practices.

Validity of Empirical Design — The key threat to identification in estimating Equations (1) and (2) is unobserved differences in the criminal environment that may be correlated with an officer's shift-hour or off-duty work schedule. With respect to Equation (1) a potential concern is that supervisors may be more likely to have officers work on patrol rather than doing some other form of work in periods in which crime is high. Therefore, certain hours may have a large share of officers working at the beginning of their shift and simultaneously experience high crime that generates higher arrest propensities. A related concern for our off-duty analysis is the possibility that officers who plan to work a second job post-shift only agree to take 911 calls and work on patrol if they expect to not see many serious calls.

Our set of fixed effects partially address these concerns by flexibly controlling for differences in arrest propensities across division $\times$ day-of-week $\times$ hour and division $\times$ day-of-week $\times$ shift, addressing all persistent differences across divisions in their criminal environments and shift-specific practices and their related variation in shift assignments and off-duty schedules. Likewise, our division $\times$ month fixed effects address changes in the environment over time. [2]

We conduct two analyses to further probe the concern that criminal activity varies throughout officers' shifts. In Figure A-16, we re-estimate our baseline analyses with the inclusion of the log volume of all calls and serious calls. The point estimates are identical to those from our main analyses. Further, as we describe later in this Section, we estimate a version of Equation (1) where the unit of observation is calls taken by officers and where

[2] In order to further probe the validity of our design, we conduct an imperfect but informative balance test for our main variables of interest. Utilizing the same regressions described in Equations (1) and (2), we place the logarithm of the total number of 911 calls and high priority 911 calls in the officer's division-hour on the left-hand side of the equation and test for whether officer shift-time predicts call volume. We present these results in Table A-11. After accounting for the set of fixed effects, the coefficient on a particular shift-hour tells us whether a division-hour with a higher than average share of officers in that shift-hour also has a higher than average number of calls. With respect to overall call volume, an F-test rejects the null hypothesis that call volume is unrelated to shift-time. However, the effects are very small in magnitude, as our shift-hour coefficients are never larger than 1 percent per hour. When we focus on high-priority calls, the p-value on the F-statistic is insignificant, indicating that there is little evidence that high-priority calls for service vary with the composition of officer shifts. The last two columns of Table A-11 test for a relationship between call volume and officer off-duty obligations. The joint F-test for column (3) also rejects the null hypothesis of insignificant coefficients, though with similarly small estimated quantities. With respect to high priority service calls, the p-value on the F-test is 0.88, indicating little evidence against balance.

we directly control for type of call, and we find a similar decline in arrest frequency near shift-end.

Functional Form — We next test the robustness of our results to choice of fixed effects. We re-estimate alternative versions of Equation (1), all of which include officer fixed effects but vary the ways in which we account for the importance of place, work duties and time. These estimates are also presented in Figure A-16. In all specifications, we find a similar pattern of arrest frequency reduction. We also observe a broadly consistent pattern for court convictions. The only exception is with the inclusion of the most granular controls that condition on division-by-date-by-hour. Including these effects means the shift-time court conviction patterns are only identified from arrests made in the exact same division, date, and hour, which are a small subset of observations.[3] The court conviction rate in later shift-hours is no longer statistically different from the first hour of the shift, but the effects are still statistically indistinguishable from our baseline patterns.

Alternative Explanations — Next, we consider whether there are alternative explanations for the decline in arrest activity that we observe towards the end of an officer's shift. One possibility is that rather than changing behavior due to an aversion to overtime work, officers reduce arrest frequency because they are fatigued from having already worked for several hours. A fatigue effect could lead to a decline in arrest propensity and an increase in the quality of realized arrests, which we would erroneously ascribe to an aversion to overtime.

To differentiate between changes in arrests due to distance from end of shift and distance from start of shift, we estimate a version of Equation 1 on an expanded sample that also includes officers who work nine-hour and ten-hour shifts. We add to our specification a variable for number of hours into the shift, which captures drivers of time-into-shift effects like fatigue. To increase precision, we include separate indicator variables for only the last four hours of a shift, so that the remaining hours are the leave-out group and contribute to identifying the time-into-shift effect.The results of this analysis are presented in Figure A-17. We plot coefficients from regressions with and without the time-into-shift variable. In both specifications, there is a clear decline in arrest propensity as officers reach the end of their shift. The robustness to including a time-into-shift effect suggests that our results are not due to a fatigue effect. In Section IV.D we further show that when officers perform off-duty work before their police shift, their arrest activity shows at most a small decline, further indicating that fatigue is not an important driver of arrest activity.

Another potential concern is that a decline in late-shift arrests could be an artifact of either a formal or informal departmental policy that routes officers to fewer calls for service towards the end of their shift. In order to address this possibility, we ask whether an officer's arrest propensity declines *conditional on taking a call for service.* We explore this question in Figure 5 which plots the regression-adjusted probability of an arrest throughout an officer's shift, where the unit of observation is a citizen 911 call for service. As in Figure 4, the relevant shift-hour coefficient is plotted on the y-axis with the hour relative to the end of the shift plotted on the x-axis. We continue to condition on all of the fixed effects in Equation

[3] Of the 90,016 observations in our baseline set of arrests, only 12,482 occur in a division-date-hour where two officers make arrests and are at different points in their shift.

(1), though, since the model is estimated at the service call level, we additionally condition on call type fixed effects. Consistent with the aggregate results, arrests decline significantly — by approximately 20 percent — throughout the officer's shift, thus indicating that our main results are not an artifact of a change in the volume of service calls at the end of the workday.

Focusing on the conditional probability of an arrest given a call for service allows us to address several additional threats to identification. First, we might be concerned that the decline in late-shift arrest activity could be the result of the incapacitative effect of early-shift arrests. That is, arrests might decline mechanically throughout the workday as officers are removed from circulation after having made an arrest in the early in their shift. By estimating the probability of an arrest conditional on a call for service, we remove the source of this concern since all of the officers taking service calls are, by definition, available to make arrests.

A second concern is that there may be measurement error in arrest timestamps. To the extent that there are systematic differences between the timing of an officer's decision to take a suspect into custody and the timestamp of the arrest in our data, we might be concerned that some arrests made later in an officer's shift might be misclassified as having occurred either earlier in the workday or after the officer's official workday has ended. Since the models in which we condition on a call for service obtain a timestamp from the service call (which is documented in the city's 911 system) rather than the arrest (which is documented by the officer), this analysis is robust to the problem of errors in the arrest data.[4]

A third concern is that some officers may engage in "arrest trading," the practice of avoiding overtime work by passing along a late-shift arrest to another officer who co-responded to a particular service call (Linn, 2009). To the extent that this practice occurs, it is possible that the decline in arrest activity that we observe at the end of an officer's shift represents a reallocation of administrative work rather than a true decline in the number of arrests that an officer makes. Figure 5 also directly addresses this issue. Because the model is estimated at the service call level, and the outcome is whether any arrest resulted, these estimates show that when an officer is dispatched to a service call at the end of his shift, an arrest is less likely — regardless of whether the arrest was made by the officer himself. This analysis rules out the possibility that arrest trading is an important contributor to our main estimates.

Finally, we have interpreted our court conviction regressions as evidence that officers' guilt threshold increases near the end of their workday. We consider here two alternative explanations. One possibility is that, for a given underlying offense, officers might increase the number of charges that they issue for a late-shift arrest, which, in the course of the plea bargaining process, could increase the probability of a conviction (Rehavi and Starr, 2014). In the left-hand panel of Figure A-18, we show how the number of arrest charges varies based on the shift-hour of the arrest. While there appears to be a small increase in the number of

[4]We can also investigate the quality of the timestamps directly. Among dispatch calls in which an arrest was made on the same day, the modal time between the dispatch call and the arrest is 0 hours, and the mean is 0.8 hours. Over 95 percent of arrests occur within 2 hours of the officers being dispatched.

charges, the magnitude is quite small and unlikely to explain our court conviction finding.

Another explanation that could potentially rationalize our findings is that officers spend more time processing late-shift arrests than they spend processing arrests that are made earlier in their shift. In that case, our conviction results could reflect changes in officer effort rather than changes to the officer's guilt threshold. To test this explanation, we consider cases when the opportunity cost of arrest processing is high, which would reduce the court conviction rate under this alternative story. First, we examine court convictions for shifts that end on Friday and Saturday between 4pm and 10pm, when leisure time may be of greatest value. In the right-hand panel of Figure A-18, we show that the time path of court convictions for these shifts, though more imprecise, are statistically indistinguishable from the full-sample time path. Second, as we discuss in Section IV.D, the presence of an off-duty job after an officer's shift increases the opportunity cost of arrest processing. In Table A-7, we show that the court conviction rate for late-shift arrests is statistically identical on days in which an officer has an off-duty job, running counter to the alternative theory that officers take their time in processing late shift arrests in order to maximize overtime pay.

D Model Details

We present here in greater detail several components of the model, including: solving the model, construction of moments, details of estimation, calculation of statistics in the model estimates table, and the equity-efficiency calculation.

Model Solution

We solve for the officer's optimal arrest behavior in each period by backwards induction. In the final period, their utility is

$$V_T(a,s) = a[(1-\lambda)\tilde{p}(s) + \phi V_{OT} + (1-\phi)V_0] + (1-a)[\lambda(1-\tilde{p}(s)) + V_0]$$

This gives us a cutoff probability $\tilde{p}_T$ above which an arrest is made, $\tilde{p}_T = \lambda + \phi[V_0 - V_{OT}]$, and an associated guilt signal s^*_T. We then calculate the probability of arrest in period T, and the probability of guilt conditional on arrest, with which we construct the value of patrol in period T:

$$\begin{aligned} V^p_T = & Pr(s > s^*_t)[(1-\lambda)Pr(g=1|s > s^*_t) + \phi V_{OT} + (1-\phi)V_0] \\ & + Pr(s < s^*_t)[\lambda(1 - Pr(g=1|s < s^*_t)) + V_0] \end{aligned}$$

If the officer is not in patrol in the final period, their utility is $V^{np}_T = \lambda(1-\pi) + pV_{OT} + (1-p)V_0$. In earlier periods, the value of not patrol is $V^{np}_t = \lambda(1-\pi) + pV^{np}_{t+1} + (1-p)V^p_{t+1}$. In periods $t < T$, utility when working and conditional on guilt probability is

$$V_t(a,s) = a[(1-\lambda)\tilde{p}(s) + \phi V^{np}_{t+1} + (1-\phi)V^p_{t+1}] + (1-a)[\lambda(1-\tilde{p}(s)) + V^p_{t+1}]$$

where we have solved for V^{np}_{t+1} and V^p_{t+1} in the previous step. Again, we solve for the cutoff probability $\tilde{p}_t = \lambda + \phi[V^p_{t+1} - V^{np}_{t+1}]$ above which an arrest is made, as well as the associated guilt signal s^*_t. From the guilt signal in time period t, we construct V^p_t.

Construction of Moments

As described in Section V, we estimate the model by matching moments from our regressions to model-generated moments. We describe each set of moments here in more detail.

Arrests — The model generated probability of arrest is a function of the joint probability that the officer is working patrol and the probability that the officer makes an arrest conditional on patrol. We denote by $\theta_t \in \{0,1\}$ the officer's patrol status in period t, where $\theta_t = 1$ indicates that the officer is on patrol. We assume that there is a probability that the officer starts on patrol, $p_{\text{start}} = Pr(\theta_1 = 1)$. If they do not start on patrol, they have a probability p of joining patrol in each period, the same probability as the processing probability.

The probability of an arrest in each period is the probability of working patrol times the probability of a guilt signal above the arrest threshold:

$$Pr(a_t = 1) = Pr(\theta_t = 1)Pr(s > s_t^*).$$

We construct the probability of patrol inductively from the officer's patrol status in the previous period,

$$Pr(\theta_t = 1) = Pr(\theta_{t-1} = 1)[Pr(s < s_{t-1}^*) + (1 - \phi)Pr(s > s_{t-1}^*)] + Pr(\theta_{t-1} = 0)(1 - p)$$

Court Conviction — The court conviction coefficient is matched to the model with the probability of guilt conditional on arrest in each period, $Pr(g = 1|s > s_t^*)$.

Arrest-Overtime Relationship — The arrest-overtime moments are the probability that the officer is *not* on patrol (i.e. processing an arrest) at the end of the day, $Pr(\theta_{T+1} = 0|a_t = 1)$, conditional on making an arrest in period t. There are two steps for this calculation. First, the probability of not patrolling in the period after arrest is set by ϕ, the probability of needing to process the arrest. Then, the probability of not being on patrol in a future period comes from either continued processing or a new arrest:

$$Pr(\theta_{t+k} = 0|a_t = 1) = \begin{cases} \phi & \text{if } k = 1, \\ Pr(\theta_{t+k-1} = 0|a_t = 1)p & \\ +Pr(\theta_{t+k-1} = 1|a_t = 1)\phi Pr(s > s_{t+k-1}^*) & \text{if } k > 1. \end{cases}$$

Notice that the arrest-overtime relationship thus captures that an officer may be working overtime because of processing a *later* arrest than the one from time period t.

Off-Duty Effects — We match four off-duty moments, which are the effects of working a short and long off-duty shift after work, separately for the first four and last four hours of the shift. For all cases, the effect is measured relative to having no off-duty obligations post-shift.

The average short off-duty shift is 3.41 hours, and the average long off-duty shift is 6.41 hours. Based on an imputed average pay rate of 35, we assign dollar values of $\text{Pay}_{OD}^{short} = 119.3$ for short off-duty shifts and $\text{Pay}_{OD}^{long} = 224.3$ for long off-duty shifts. From these, we construct the value of ending the work day with overtime as $V_0 = 0$ for no off-duty days, and $V_{OD}^{short} = c_{OD} + b\text{Pay}_{OD}^{short}$ and $V_{OD}^{long} = c_{OD} + b\text{Pay}_{OD}^{long}$ for short and long off-duty days, respectively. We add here the notation $OD = \{0, s, l\}$ to indicate whether the officer has no, short, or long off-duty at the end of the day.

We then calculate the arrest probability for each hour, a_t, separately by all three off-duty possibilities. Our model-simulated moments are then:

$$\Delta^{short}_{early} = \sum_{t=1}^{4} E(a_t|OD = s) - E(a_t|OD = 0), \ \Delta^{short}_{late} = \sum_{t=5}^{8} E(a_t|OD = s) - E(a_t|OD = 0)$$
$$\Delta^{long}_{early} = \sum_{t=1}^{4} E(a_t|OD = l) - E(a_t|OD = 0), \ \Delta^{long}_{late} = \sum_{t=5}^{8} E(a_t|OD = l) - E(a_t|OD = 0)$$

Model Estimation — We estimate the model by choosing the set of parameters θ that minimize the distance between the model-based moments $M_k(\theta)$ and empirical moments, $\hat{M}_k$,

$$\hat{\theta} = \arg\min_{\theta} \sum_k [M_k(\theta) - \hat{M}_k]^2/\hat{s}_k$$

where $\hat{s}_k$ is the estimated standard error for empirical moment k. We solve for the minimum using the fmincon function in Matlab. To avoid the possibility of finding a local but non-global optimum, we estimate the model multiple times with a range of starting parameter guesses and choose the parameter estimates that minimizes the objective function across all estimations.

Intuition for Identification from Short-versus-Long Off-Duty

We describe here the intuition for how we are using the after-shift off-duty impacts to identify our model parameters and the crucial importance of variation in length of off-duty shifts.

Suppose that an officer in their last shift-hour faces an individual with probability $\tilde{p}$ of guilt. For simplicity, we suppose that an arrest will with certainty require processing, $\phi = 1$. If the officer makes an arrest, they receive utility $(1-\lambda)\tilde{p} + V_{OT} = (1-\lambda)\tilde{p} + c_{OT} + b\text{Pay}_{OT}$. If they do not make an arrest, they receive utility $\lambda(1-\tilde{p}) + V_0 = \lambda(1-\tilde{p}) + \mathbb{I}[OD = 1] \cdot (c_{OD} + b \cdot \text{Pay}_{OD})$, where V_0 reflects the officer's value of ending their shift without overtime work: if they do not have off-duty work, they receive the normalized value of going home, 0. If they have off-duty work, they receive $c_{OD} + b \cdot \text{Pay}_{OD}$.

Therefore, the officer makes an arrest if $(1-\lambda)\tilde{p} + c_{OT} + b\text{Pay}_{OT} \geq \lambda(1-\tilde{p}) + \mathbb{I}[OD = 1] \cdot (c_{OD} + b \cdot \text{Pay}_{OD})$, which translates to the condition

$$\tilde{p} \geq \lambda + \mathbb{I}[OD = 1] \cdot (c_{OD} + b \cdot \text{Pay}_{OD}) - (c_{OT} + b\text{Pay}_{OT})$$

This equation shows that the presence of after-shift off-duty work changes the officer's guilt threshold by $\mathbb{I}[OD = 1] \cdot (c_{OD} + b \cdot \text{Pay}_{OD})$, which does not alone allow us to separately identify the parameters c_{OD} and b. However, variation in the *length* of the officer's off-duty shift induces differences in Pay_{OD}, which only shifts the officer's guilt threshold by $b \cdot \text{Pay}_{OD}$ and allows us to identify b and thus c_{OD} (and c_{OT}).

Table 4 Estimate Description

Overtime pay needed for arrest of certainly innocent suspect — We solve for this figure based on the value of Pay_{OT} needed such that, when $\tilde{p} = 0$, an officer is indifferent between arrest and non-arrest in the final hour of the day. The value of patrol in the final hour in this case is $V_T(a, s = -\infty) = (1-a)\lambda + a\phi[c_{OT} + b \cdot Pay_{OT}]$. We solve for the value of Pay_{OT} where $V_T(0, -\infty) = V_T(1, -\infty)$:

$$\begin{aligned} \lambda &= \phi[c_{OT} + b \cdot Pay_{OT}] \\ Pay_{OT} &= (\lambda/\phi - c_{OT})/b \end{aligned}$$

Off-duty pay needed for non-arrest of certainly guilty suspects — We solve for this figure based on the value of Pay_{OD} needed such that, when $\tilde{p} = 1$, an officer is indifferent between arrest and non-arrest in the final of the hour of the day when an officer has off-duty work scheduled after their shift. The value of patrol in the final hour in this case is $V_T(a, +\infty) = a(1-\lambda) + a[\phi(c_{OT} + b \cdot Pay_{OT}) + (1-\phi)(c_{OD} + b \cdot Pay_{OD})] + (1-a)(c_{OD} + b \cdot Pay_{OD})$. We solve for the value of Pay_{OD} where $V_T(0, +\infty) = V_T(1, +\infty)$:

$$\begin{aligned} (1-\lambda) + [\phi(c_{OT} + b \cdot Pay_{OT}) + (1-\phi)(c_{OD} + b \cdot Pay_{OD})] &= (c_{OD} + b \cdot Pay_{OD}) \\ Pay_{OD} &= [(1-\lambda)/\phi + (c_{OT} - c_{OD})]/b + Pay_{OT} \end{aligned}$$

Equity-Efficiency Calculation for Production Possibilities Frontier

To calculate the equity-efficiency frontier in Figure A-8, we first remove overtime considerations by setting $V_{OT} = V_{OD} = 0$, so officers are indifferent between ending the day on or off patrol. By doing so, officers' objective function is solely to maximize the weighted average of arrests of guilty suspects and non-arrests of innocent suspects. We then span every value $\lambda \in [0, 1]$, and we calculate the associated values of $Pr[a = 1|g = 1]$ and $Pr[a = 0|g = 0]$:

$$\begin{aligned} Pr[a = 1|g = 1] &= \frac{1}{8}\sum_{t=1}^{8} Pr[a_t = 1|g_t = 1] = \frac{1}{8}\sum_{t=1}^{8} Pr(\theta_t = 1) \cdot Pr(s > s_t^*|g_t = 1) \\ Pr[a = 0|g = 0] &= \frac{1}{8}\sum_{t=1}^{8} Pr[a_t = 0|g_t = 0] = \frac{1}{8}\sum_{t=1}^{8} [1 - Pr(\theta_t = 1) \cdot Pr(s > s_t^*|g_t = 0)] \end{aligned}$$

where s_t^* and $Pr(\theta_t = 1)$ depend on λ. These calculations give us the equity-efficiency frontier. The "true value" point in the figure is calculated using the estimated values of λ and V_{OT}. The "overtime indifferent" point in the figure is calculated using the estimated value of λ but setting $V_{OT} = 0$.

E Accounting for Career Incentives

We argue that, based on institutional details, officer arrest activity does not play a role in department promotion decisions or in the receipt of off-duty work. In this section, we demonstrate empirically that an officer's promotion probability and volume of off-duty work are not related to their volume of arrests. We then develop an extended version of our structural model that explicitly allows for an economic incentive for arrests.

Arrests and Promotion — We measure officer rank from two snapshot rosters of the Dallas Police Department from October 2017 and September 2021.[5] We examine promotions to Corporal and Sergeant, which are the two ranks above Officer. For each officer, we create a count of all arrests made from January 2018 to March 2019 (the end of our sample), separately by whether the arrest leads to a guilty conviction in court (we'll refer to these as "guilty arrests" and "not-guilty arrests"). For the Corporal (Sergeant) analysis, we restrict attention to employees who are at Officer (Corporal) rank in 2017. We also restrict attention to employees whose tenure is long enough to be eligible for promotion.

Our specification will be a cross-sectional regression,

$$\text{Promoted}_i = \alpha_g \text{GuiltyArrests}_i + \alpha_{ng} \text{NotGuiltyArrests}_i + X_i \beta + \epsilon_i$$

where X_i includes a set of additive fixed effects for race, gender, and education, as well as fixed effects for modal division in which an officer worked in the sample period, and fixed effects for months of tenure.

As part of this analysis, we are also interested in the impact of promotion on salary. We will next place promotion as a right-hand side variable, and we place 2021 salary as the dependent variable. We similarly include tenure fixed effects but omit demographic and division fixed effects.

The results of this analysis are presented in Table A-10. Column 1 and 3 (2 and 4) focus on the sample of individuals who have not been promoted to Corporal (Sergeant) by October 2017 but have sufficiently long tenure to potentially be promoted by September 2021. The estimates in Column 1 show that officers who make more arrests of guilty individuals are *not* more likely to be promoted to corporal than similarly-situated officers who made fewer guilty arrests. Similarly, there is no difference in promotion rate based on number of arrests of non-guilty individuals. The patterns we find in Column 2 indicate that promotion to sergeant is similarly unrelated to volume of guilty and non-guilty arrests.

Notice that, since these regressions are based on cross-sectional correlations in arrest volume and promotion, a significant relationship may have been possible even absent a causal effect of arrest activity on promotion probability. Nevertheless, the null findings provide reassurance that the institutional rules are followed in practice and arrests are not factored into promotion decisions.

[5]These rosters also include individuals who have left the department at some point since 2010, so we are able to include individuals who are promoted between October 2017 to September 2021 but leave the department before September 2021.

Columns 3 and 4 show that promotions have a sizable impact on officer salary. Relative to an officer with the same tenure, an officer who is promoted to Corporal (Sergeant) receives on average \$5261.80 (\$6846.80) more in annual salary.

While the arrest coefficients are insignificant, it is useful to consider the magnitude of the implied pecuniary benefit of arrests. We take the coefficients in Columns 1 and 2 and multiply them by the coefficients in Columns 3 and 4. We then multiply these estimates by the average number of years officers remain on the force conditional on being present at the eligibility tenure (which is roughly 20 years for both ranks). Note that these calculations overestimate the economic returns to a promotion, since they do not account for the possibility that officers may be promoted at a future date.

For corporal, the regression estimates imply an "economic return" of \$394.0 to a guilty arrest and \$42.4 to a not-guilty arrest. For sergeants, the implied return is -\$386.8 for a guilty arrest and \$49.7 for a not-guilty arrest. In the next analysis, where we incorporate career incentives into our model, we will use these estimated magnitudes as a guide for reasonable incentives to consider.

Arrests and Off-Duty Opportunities — We will next assess the impact of arrest activity on the takeup of off-duty work. We construct an officer-month panel, where we calculate a moving summation for number of arrests in the previous three months, separately by guilty and not-guilty. We regress the number of off-duty hours taken by an officer in a given month on the last three months of arrests:

$$\text{OffDutyHours}_i = \alpha_g \text{GuiltyArrests}_i + \alpha_{ng} \text{NotGuiltyArrests}_i + \theta_t + \mu_{d(i,t)} + \epsilon_i$$

where θ_t is a set of year-month fixed effects, and $\mu_{d(i,t)}$ is a set of fixed effects for the division an officer is assigned to in that month. We present specifications with and without officer fixed effects.

The results of this analysis are presented in Table A-12. The first column does not include officer fixed effects. We find a small and statistically insignificant coefficient on volume of guilty arrests. The coefficient on volume of non-guilty arrests is statistically significant but also small. In the second column, we add officer fixed effects. The coefficients for both variables decline in magnitude, and both are now statistically insignificant.

To provide a sense of magnitude, officers make approximately \$35 per hour in the typical off-duty position.[6] The coefficients in the first column indicate that a guilty arrest is associated with $-0.208 \times 35 \times 3 = -\$21.84$ in off-duty work, where the multiplication by three is because an arrest appears in the independent variable for three separate months. Similarly, a not-guilty arrest is associated with $-0.233 \times 35 \times 3 = -\$24.46$ in off-duty work. These values are much smaller than the insignificant estimates for promotion incentives, and they are consistent with the institutional details that state that officers are allowed to assume off-duty work, regardless of their workplace arrest activity. As we show below, much larger economic returns to arrests are necessary to change our estimates of professional motivations.

[6]As we highlight in the paper, we do not see off-duty salaries in our data, and this number is based on asking officers and information online.

Model Extension with Promotion Concerns — In this section, we extend our model from Section V to explicitly allow that arrests might lead to a personal gain to officers. Equation (3) shows the officer's per-period utility from their arrest activity as a function of their arrest choice and their belief about suspect guilt. We modify this expression to incorporate an impact on promotion outcomes:

$$v^p(a,s) = \underbrace{a(1-\lambda)\tilde{p}(s) + (1-a)\lambda(1-\tilde{p}(s))}_{\text{Professional Motivations}} + \underbrace{a[\tilde{p}(s)b\Delta_g + (1-\tilde{p}(s))b\Delta_{ng}]}_{\text{Career Incentives}}$$

The "Professional Motivations" expressions are the same as before. The new "Career Incentives" expression states that, if an officer makes an arrest, there is probability $\tilde{p}$ that the suspect is guilty, in which case the officer receives utility $b\Delta_g$, which is the monetary benefit in terms of increased promotion probability, scaled by the dollar utility coefficient. With probability $(1-\tilde{p})$, the suspect is innocent, and the monetary benefit (or cost) in terms of promotion probability is $b\Delta_{ng}$.

With this modified per-period utility function, the officer has a new optimal cutoff guilt probability, above which he makes an arrest:

$$\tilde{p}(s_t^*) = \frac{\lambda - b\Delta_{ng} + \phi\left[V_{t+1}^p - V_{t+1}^{np}\right]}{1 + b(\Delta_g - \Delta_{ng})}$$

Notice that, when $\Delta_g = \Delta_{ng} = 0$, this cutoff rule reduces to the previous rule without career incentives.

We solve for the new optimal strategy for the officer with the same backwards induction approach as before: we solve for the optimal cutoff in period $T = 8$, which generates the expected values for being on and off patrol in that period. We then use those values to solve for the optimal cutoff in period $t = 7$, and so on.

The estimation uses the same set of moments as our baseline model. While our regression analysis above indicates that $\Delta_g \approx \Delta_{ng} \approx 0$, we will estimate the model with several different candidate values for Δ_g and Δ_{ng}. We will consider the following four values for career incentives: 1) guilty arrests are rewarded ($\Delta_g = 1000$, $\Delta_{ng} = 0$); 2) innocent arrests are punished ($\Delta_g = 0$, $\Delta_{ng} = -1000$); 3) both arrests are rewarded ($\Delta_g = \Delta_{ng} = 1000$); and guilty arrests are rewarded and innocent arrests are punished ($\Delta_g = 1000$, $\Delta_{ng} = -1000$). We choose a magnitude of \$1000 because it is roughly 2-3x greater than the largest (statistically insignificant) magnitudes found across Tables A-10 and A-12.

The results of this new estimation are presented in Table A-9 below. The first column presents the baseline estimates without career incentives, and the next four columns present estimates with varying types of career incentives. Column 2 shows that, perhaps surprisingly, incorporating a career incentive to make guilty arrests actually *increases* the implied professional motivation towards non-arrest of innocent suspects. Notice that the increase is due to a decline in b. The likely driver of this result is that, as officers face greater career incentives to make arrests, the small magnitude of the off-duty arrest decline becomes more consistent with a larger professional motivation. This logic also applies in Column 4 — when both guilty and innocent arrests are rewarded — where we similarly see large increases in both measures of professional motivation.

In Column 3, where wrongful arrests are punished, the professional motivation towards innocent non-arrest is essentially unchanged, and the motivation towards guilty arrests declines somewhat, from \$4042 to \$2456. Here again the most important driver of the change seems to be a change in b, which now increases.

In Column 5, we consider the case where guilty arrests are rewarded and non-guilty arrests are punished. Here, the value of non-guilty non-arrest increases from \$616 to \$1,941, and the value of a guilty arrest decreases from \$4042 to \$1047.

These alternative specifications demonstrate the importance of carefully considering exactly the types of economic incentives bureaucrats face, since the particular arrangement can have important impacts on estimates of professional motivation. However, these results provide reassurance that our estimates are not biased upwards by the possible presence of career concerns.

F Accounting for Officer Fatigue

We argue that the observed decline in arrest propensity near shift-end is due to officers' aversion to working overtime. One competing hypothesis is that this pattern is due to officer fatigue. In this section, we explicitly consider the possibility of officer fatigue in our model. We do so by incorporating a private cost of making an arrest which increases with time into shift.

We modify our model by adding an expression to officers' per-period utility,

$$v_t^p(a,s) = \underbrace{a(1-\lambda)\tilde{p}(s) + (1-a)\lambda(1-\tilde{p}(s))}_{\text{Professional Motivations}} - \underbrace{a\kappa(t-1)}_{\text{Fatigue}},$$

where the expression is such that the fatigue cost is 0 in period $t = 1$ and increases linearly throughout the shift. This additional term leads to a per-period cutoff rule:

$$\tilde{p}(s_t^*) = \lambda + \kappa(t-1) + \phi[V_{t+1}^p - V_{t+1}^{np}]$$

When $\kappa = 0$, the guilt threshold for arrests only differs across periods because of arrests' impact on officer incapacitation and overtime, $\phi[V_{t+1}^p - V_{t+1}^{np}]$. With $\kappa > 0$, fatigue induces an increase in the guilt threshold throughout the shift, separate from incapacitation/overtime considerations.

To estimate this additional parameter, we exploit variation in whether officers worked an off-duty shift *prior* to their shift. If an officer worked H_{OD} off-duty hours prior to their police shift, we assume that the fatigue expression takes the form $a\kappa(t - 1 + H_{OD})$. In our data, the average pre-shift off-duty work shift is 5.23 hours. Denoting by $OD = \{0, b\}$ whether an officer had no off-duty or a before-work off-duty shift, we simulate two additional moments that we match to the data: the impact of a before-work off-duty shift on early and late arrests:

$$\Delta_{early}^{before} = \sum_{t=1}^{4} E(a_t|OD = b) - E(a_t|OD = 0)$$

$$\Delta_{late}^{before} = \sum_{t=5}^{8} E(a_t|OD = b) - E(a_t|OD = 0)$$

For our empirical moments, we borrow our estimates from Column 1 in Table 3, which estimates the impacts of regular off-duty shifts before and after an officer's shift separately by early and late in the shift. We estimate that a before-work off-duty shift induces a 0.000120 (0.000348) impact on early-shift arrests and a -0.000152 (0.000344) impact on late-shift arrests, neither of which is statistically significant.

We re-estimate all the model's parameters, plus the inclusion of κ. We estimate a value of $\kappa = 0.00018$. This magnitude is very small, and indicates that fatigue causes officers' guilt threshold to increase by $7 \times 0.00018 = 0.0013$ (where the scale is 0 to 1) in the last shift-hour relative to the first shift-hour. If fatigue effects were a meaningful driver of the arrest decline, we would expect a sizable change in c_{ot}/b, which measures the non-monetary cost of working overtime, denominated in dollars. This figure is -1081.0 in our baseline estimation and is -1054.0 in our specification with officer fatigue.

Another way to understand the magnitude of the fatigue effect is to simulate the officers' arrest probability with our new model estimates, where we counterfactually set $\kappa = 0$ and leave all other parameters unchanged. This calculation is presented in Figure A-19. The red line presents the model fit to the empirical arrest pattern. The dotted line presents the counterfactual where $\kappa = 0$. The time path of arrests is nearly identical, which strongly indicates that the shift-end decline in arrests is not due to fatigue.

Interpreting Arrest Differences by Shift-Length

In Appendix Section C and Figure A-17, we present evidence that the arrest decline near shift-end is not due to time into shift, and we do so by expanding our sample to include 9-hour and 10-hour shifts.

Through the lens of this model formulation, what should we expect the 9-10 hour shifts to do to officer arrests? We will modify our notation slightly, and $t = -10, ..., -1$ now denotes hours from shift-end, where officers with different shift lengths start with different values of t. In addition, $\tau = 1, ..., 10$ denotes hours *into* shift. The optimal cutoff rule has the following form:

$$\tilde{p}(s^*_{\tau,t}) = \lambda + \underbrace{\kappa(\tau - 1)}_{\substack{\text{Function of} \\ \text{time from shift start}}} + \underbrace{\phi[V^p_{t+1} - V^{np}_{t+1}]}_{\substack{\text{Function of} \\ \text{time from shift end}}}$$

There are two insights here. First, the presence of shifts with different lengths allows us to separately identify time-into-shift effects from time-from-shift-end effects. Second, while fatigue has a slope effect throughout the shift, the difference in arrest activity across shift lengths is a *level* effect: relative to an 8-hour shift, $\tilde{p}(s^*_{\tau,t})$ is κ higher for 9-hour shifts and 2κ higher for 10-hour shifts. These insights accord with our choice reduced-form regression specification, where there is a single time-into-shift linear term across all shift length and a single set of time-from-shift-end indicators.

One small detail here is that we do not observe $\tilde{p}(s^*_{\tau,t})$ directly. The regressions are instead on arrest probability, $Pr(\text{arrest}) = Pr(s > s^* \;\; \& \;\; \text{officer on patrol})$. So, in the reduced-form regressions, the choice of a linear time-into-shift effect is an approximation to a likely non-linear effect.

Table A-1: Summary Statistics for Ten Most Common Arrest Types

UCR Offense	(1) Share of Arrests	(2) Felony	(3) Convicted	(4) Hour of Shift	(5) Overtime Paid
All Arrests	1.00	0.21	0.34	4.22	0.45
Warrant	0.28	0.00	0.06	4.24	0.39
Assault	0.15	0.27	0.29	4.18	0.49
Disorderly Conduct	0.14	0.00	0.08	4.18	0.34
Narcotics & Drugs	0.13	0.51	0.71	4.30	0.49
Public Intoxication	0.03	0.01	0.07	4.16	0.37
Dwi	0.03	0.11	0.88	4.26	0.61
Trespass	0.02	0.02	0.81	3.99	0.44
Theft-Retail	0.02	0.32	0.74	4.13	0.45
Other-Misdemeanor	0.02	0.00	0.12	3.67	0.35
Theft-Other	0.02	0.70	0.73	4.16	0.56

Note: Table presents summary statistics for the ten most common types of arrest charges in the data. For each arrest charge, we note the share of arrests, the share that are felonies, the share resulting in a conviction, the mean shift-hour and the probability that the arrest charge leads to an overtime spell.

Table A-2: "First Stage" Regression: Relationship Between Predicted and Actual Off-Duty Work

	(1) Before OD	(2) After OD
Regular Off-Duty Before Shift	0.509	0.0292
	(0.00586)	(0.00291)
Regular Off-Duty After Shift	-0.000733	0.516
	(0.00150)	(0.00555)
Mean	0.063	0.095
Observations	614920	614920

Standard errors in parentheses

Note: Table presents estimates from a regression of an indicator for off-duty work on an indicator for predicted off-duty work, where the predicted work variable is equal to 1 if the officer worked an off-duty shift on a given day of the week at least 25 percent of the time in a given quarter. Models condition on the fixed effects described in Equation (1). Standard errors are clustered at the division-by-month level.

Table A-3: Most Common Off-Duty Jobs

American Airlines Center (AAC)
Greenway Parks Home Owners Association
Kalua Discoteque
Green Oaks Hospital
Cowboys Red River
Prestonwood PID ENP
Crescent Hotel
Hunt Oil
Southwest Center Mall
Preston Hollow North Inc ENP
North Bluffview ENP
Texas Scottish Rite Hospital for Children
Children's Medical Center
Bank of America Plaza
Inwood National Bank
Pegasus Link Constructors, LLC
Royalwood ENP
Watermark Church
Meadows Foundation Inc
Medical City Dallas Hospital ER

Note: Table presents a list of the most common off-duty jobs worked by Dallas police officers during the study period.

Table A-4: Robustness of Off-Duty Estimates to Alternative Models

	(1)	(2)	(3)	(4)	(5)	(6)
	Arrest	Arrest	Arrest	Arrest	Arrest	Arrest
Regular Off-Duty	0.000120	0.000150	0.000221	0.000147	0.000308	0.000148
Before × Early	(0.000348)	(0.000346)	(0.000342)	(0.000345)	(0.000443)	(0.000450)
Regular Off-Duty	-0.000152	-0.000100	-0.000143	-0.000123	0.0000916	-0.000100
Before × Late	(0.000344)	(0.000349)	(0.000336)	(0.000343)	(0.000399)	(0.000407)
Regular Off-Duty	-0.000132	-0.000223	-0.000218	-0.000186	-0.000220	-0.000156
After × Early	(0.000320)	(0.000325)	(0.000311)	(0.000320)	(0.000390)	(0.000393)
Regular Off-Duty	-0.000521	-0.000607	-0.000549	-0.000497	-0.000726	-0.000671
After × Late	(0.000262)	(0.000264)	(0.000255)	(0.000263)	(0.000310)	(0.000313)
Mean	0.017	0.017	0.017	0.017	0.017	0.017
Observations	4919608	4919608	4919595	4917539	4909911	4909911
Baseline FE	X					
Baseline + Division-Date FE		X				
Division-Month-Weekend-Hour FE			X			
Division-Date-Hour FE				X		
Badge-Quarter-Hour					X	X
Division-DOW-Hour						X

Standard errors in parentheses

Note: Table presents estimates of variants on Equation (2) with different sets of fixed effects. Standard errors are clustered at the division-by-month level.

Table A-5: Effect of Off-Duty Employment, Short versus Long Off-Duty Shift Length

	Hourly Sample		Dispatch Sample	
	(1)	(2)	(3)	(4)
	Arrest	Arrest	Arrest	Arrest
Regular Off-Duty Before × Early × Short	0.000239	0.000238	-0.000133	-0.0000612
	(0.000441)	(0.000442)	(0.000908)	(0.000908)
Regular Off-Duty Before × Late × Short	-0.000218	-0.000212	-0.00270	-0.00281
	(0.000423)	(0.000423)	(0.00105)	(0.00105)
Regular Off-Duty Before × Early × Long	-0.0000820	-0.0000892	0.00154	0.00155
	(0.000508)	(0.000508)	(0.00113)	(0.00113)
Regular Off-Duty Before × Late × Long	-0.0000495	-0.0000369	-0.000664	-0.000594
	(0.000501)	(0.000501)	(0.00138)	(0.00137)
Regular Off-Duty After × Early × Short	-0.0000244	-0.0000216	-0.000967	-0.000939
	(0.000391)	(0.000391)	(0.000849)	(0.000849)
Regular Off-Duty After × Late × Short	-0.000263	-0.000259	-0.00234	-0.00233
	(0.000308)	(0.000308)	(0.000857)	(0.000855)
Regular Off-Duty After × Early × Long	-0.000308	-0.000297	-0.000745	-0.000710
	(0.000481)	(0.000481)	(0.000892)	(0.000893)
Regular Off-Duty After × Late × Long	-0.000912	-0.000903	-0.00241	-0.00238
	(0.000397)	(0.000398)	(0.00112)	(0.00112)
Log Calls		-0.00198		-0.0131
		(0.000231)		(0.000564)
Log Serious Calls		0.00289		-0.00272
		(0.000181)		(0.000454)
Mean	0.017	0.017	0.032	0.032
Observations	4919608	4919608	1754414	1754414

Standard errors in parentheses

Note: Table presents estimates from a series of regressions of an arrest indicator on the interaction between either predicted before or after-shift off-duty work and indicators for 1) whether a shift-hour is early (first four hours) or late (second four hours) in an officer's shift and 2) whether the off-duty is long (more than 4 hours) or short (less than four hours). Standard errors are clustered at the division-by-month level.

Table A-6: Effects of Off-Duty Employment, Close versus Far from Police Shift

	(1)	(2)	(3)
	Arrest	Arrest	Arrest
Regular After	0.000442	-0.000247	-0.000247
× Early × Close	(0.000281)	(0.000360)	(0.000360)
Regular After	-0.000416	-0.000569	-0.000569
× Late × Close	(0.000260)	(0.000319)	(0.000319)
Regular After	0.0000787	-0.000302	-0.000302
× Early × Far	(0.000496)	(0.000620)	(0.000620)
Regular After	0.000492	0.000122	0.000122
× Late × Far	(0.000362)	(0.000433)	(0.000433)
Mean	0.017	0.017	0.017
Observations	4919608	4909911	4909911
Baseline FE	X		
Badge-Quarter-ShiftHour FE		X	X
Division-Week-Hour FE			X

Standard errors in parentheses

Note: Table presents estimates from a series of regressions of an arrest indicator on the interaction between regular after-shift off-duty work and indicators for 1) whether a shift-hour is early (first four hours) or late (second four hours) in an officer's shift and 2) whether the off-duty is is directly after ("close") or long after ("far") the officer's work shift ends. Each model conditions on a different set of fixed effects. Standard errors are clustered at the officer and division-by-month level.

Table A-7: Effect of Off-Duty Employment on Arrest Outcomes

	(1) Guilty	(2) AnySentence	(3) Felony	(4) Mis/Fel
Regular Off-Duty	-0.00872	0.00319	-0.00773	-0.00299
Before × Early × Short	(0.0116)	(0.00944)	(0.0103)	(0.0122)
Regular Off-Duty	-0.00121	0.00406	0.000200	-0.00959
Before × Late × Short	(0.0130)	(0.0102)	(0.0116)	(0.0131)
Regular Off-Duty	0.00690	0.0143	0.00961	-0.00251
Before × Early × Long	(0.0157)	(0.0134)	(0.0133)	(0.0164)
Regular Off-Duty	0.0102	0.0106	0.0332	0.0312
Before × Late × Long	(0.0166)	(0.0145)	(0.0157)	(0.0170)
Regular Off-Duty	-0.0188	-0.00395	-0.00202	-0.00716
After × Early × Short	(0.0114)	(0.00953)	(0.00978)	(0.0118)
Regular Off-Duty	-0.00564	-0.00156	-0.0132	-0.0202
After × Late × Short	(0.0129)	(0.0112)	(0.0118)	(0.0139)
Regular Off-Duty	-0.00392	0.00571	0.00100	0.0188
After × Early × Long	(0.0134)	(0.0109)	(0.0117)	(0.0134)
Regular Off-Duty	0.00220	0.0220	-0.0147	-0.0118
After × Late × Long	(0.0168)	(0.0143)	(0.0137)	(0.0167)
Log Calls	-0.00641	0.00470	-0.0220	-0.0137
	(0.00590)	(0.00493)	(0.00505)	(0.00625)
Log Serious Calls	0.000897	-0.00786	0.0231	0.0134
	(0.00540)	(0.00446)	(0.00443)	(0.00554)
Mean	0.327	0.186	0.194	0.455
Observations	89834	89834	89834	89834

Standard errors in parentheses

Note: Table presents estimates from a series of regressions of different outcomes of an arrest on an indicator for whether the officer has a regular off-duty shift either before or after their police shift. Models condition on the fixed effects described in Equation (2). Standard errors are clustered at the division-by-month level.

Table A-8: Model Parameter Estimates, Alternative Specifications

Parameters	Baseline	Court Errors	Biggest OD Decline	Both	Description
ρ	0.032	0.019	0.032	0.018	Probability of guilt
μ	1.595	2.033	1.617	2.067	Mean of signal for guilty individuals (officer ability)
λ	0.132	0.063	0.129	0.059	Tradeoff for arresting guilty and not arresting innocent individuals
c_{ot}	-0.227	-0.236	-0.290	-0.304	Intercept value/cost of working overtime
c_{od}	-0.018	-0.019	-0.061	-0.066	Intercept value/cost of working off-duty
b	0.00021	0.00022	0.00058	0.00062	Value of $1 of overtime/off-duty pay
ϕ	0.741	0.742	0.740	0.741	Probability or processing an arrest
p	0.765	0.764	0.766	0.765	Per-period probability of continued arrest processing
p_{start}	0.349	0.336	0.348	0.335	Probability of beginning work on patrol
p_{ot}	0.290	0.290	0.290	0.290	Probability of receiving overtime when not making an arrest
Parameters					**Description**
λ/b	615.9	280.4	224.4	95.5	Dollar value of non-arrest of innocent pers
$(1-\lambda)/b$	4042.1	4184.9	1510.7	1515.0	Value of arrest of guilty person
$E(v(a,s)\|a=0)$	601.9	277.6	219.5	94.6	Average value of non-arrest
$E(v(a,s)\|a=1)$	1407.9	1301.1	526.9	472.1	Average value of arrest
$E(\max_{a_t} v(a,s)) - v_{np}$	29.7	33.1	11.2	12.2	Average value of hour on patrol (relative to hour not on patrol)

Note: Table presents parameter estimates from the model in Section V under alternative modeling choices descri in Section V.E.

Table A-9: Model Parameter Estimates with Career Concerns

	Baseline	Guilty Arrest Incentives	Innocent Arrest Penalty	High Both Incentives	Opposing Incentives	Description
Career Incentives						
Δ_g	0	1000	0	1000	1000	Career value of guilty arrest
Δ_{ng}	0	0	-1000	1000	-1000	Career value of innocent arrest
Parameter Estimates						
ρ	0.032	0.037	0.038	0.033	0.040	Guilt prob.
μ	1.595	1.476	1.452	1.569	1.427	Guilty signal mean (officer ability)
λ	0.132	0.428	0.225	0.270	0.649	Officer risk preference
c_{ot}	-0.227	-0.235	-0.254	-0.210	-0.281	Intercept value of overtime
c_{od}	-0.018	-0.018	-0.031	-0.005	-0.031	Intercept value of off-duty
b	0.00021	0.00023	0.00032	0.00013	0.00033	Value of $1 of overtime/off-duty pay
ϕ	0.741	0.742	0.742	0.745	0.747	Arrest processing prob.
p	0.765	0.765	0.764	0.765	0.764	Continued arrest processing prob.
p_{start}	0.349	0.338	0.334	0.349	0.320	Prob. of beginning work on patrol
p_{ot}	0.290	0.290	0.290	0.290	0.290	Prob. of receiving overtime when not making an arrest
λ/b	615.9	1880.8	712.9	2010.4	1940.7	$ value of innocent non-arrest
$(1-\lambda)/b$	4042.1	2518.1	2456.2	5446.5	1047.4	Value of guilty arrest
$E(v(a,s)\|a=0)$	601.9	1828.9	692.4	1962.8	1881.6	Average non-arrest value
$E(v(a,s)\|a=1)$	1407.9	1223.3	202.3	2894.0	65.4	Average arrest value
$E(\max_{a_t} v(a,s)) - v_{np}$	29.7	-0.1	-7.9	46.6	-36.5	Average value of patrol hour (relative to non-patrol hour)

Note: This table presents estimates of our model parameters under different assumptions about officer economic incentives for arrests, as described in Section E.

Table A-10: Cross-Sectional Relationship of Arrests, Promotion, and Salaries

	(1) Corporal Promotion	(2) Sergeant Promotion	(3) Annual Salary	(4) Annual Salary
Number of Guilty Arrests	0.00382	-0.00305		
	(0.00249)	(0.00315)		
Number of Not Guilty Arrests	0.000412	0.000414		
	(0.00178)	(0.00223)		
Promoted to Corporal			5261.8	
			(606.8)	
Promoted to Sergeant				6846.8
				(768.0)
Mean	0.302	0.237	80211	89633
Observations	681	372	757	445

Standard errors in parentheses

Note: This table presents a set of regressions on the relationship between officer arrest volume, promotion, and annual salary. For each officer, we calculate the total number of arrests for the period January 2018 to March 2019 (end of sample), separately by whether the arrest leads to a guilty disposition in court. Promotion and salary information are taken from a September 2021 roster. Columns 1 and 3 restrict attention to officers who are lower-ranked than Corporal in October 2017 (our earlier roster) and who have sufficient tenure to be eligible for promotion by September 2021. Columns 2 and 4 restrict attention to officers who are lower-ranked than Sergeant in October 2017 and who have sufficient tenure to be eligible for promotion by September 2021. Columns 2 and 4 further restrict attention to officers who are at least a corporal rank and thus eligible for promotion to sergeant. The outcome of Column 1 (2) is an indicator for the officer having a corporal (sergeant) rank in October 2017. The outcome of Columns 3 and 4 are annual salary. In all four regression specifications, we include a set of tenure-month fixed effects, as well as a set of fixed effects for the modal division in which each officer worked in January 2015 to June 2017. Robust standard errors are used.

Table A-11: Balance Table: 911 Calls for Service by Shift-Hour

	(1) Log Calls	(2) Log Serious Calls	(3) Log Calls	(4) Log Serious Calls
1hr From End	-0.00416	-0.00242		
	(0.00231)	(0.00226)		
2hr From End	-0.00413	-0.00251		
	(0.00253)	(0.00214)		
3hr From End	-0.00282	0.00190		
	(0.00243)	(0.00214)		
4hr From End	-0.000938	0.00441		
	(0.00245)	(0.00209)		
5hr From End	-0.000454	0.00415		
	(0.00241)	(0.00217)		
6hr From End	-0.000280	0.00000167		
	(0.00187)	(0.00174)		
7hr From End	0.000250	-0.0000299		
	(0.00145)	(0.00158)		
Before-Shift Off-Duty			-0.00124	-0.00175
			(0.000995)	(0.00108)
After-Shift Off-Duty			0.00332	0.000199
			(0.00114)	(0.00102)
Mean	2.040	1.452	2.040	1.452
Observations	4919608	4919608	4919608	4919608
F-value	0.773	3.467	5.149	1.342
F-test	0.610	0.001	0.006	0.263

Standard errors in parentheses

Note: Table presents estimates from a regression of the natural logarithm of the total number of 911 calls and 911 calls for high-priority calls on a vector of shift-hour dummies, conditional on the fixed effects in Equation (1). Results are presented separately for the hourly sample and the dispatch sample. Below the estimated coefficients and standard errors we report the F-statistic along with the its associated p-value on the joint significance of the shift-hour terms in predicting the number of service calls.

Table A-12: Arrest Activity and Off-Duty Hours

	(1) Off-Duty Hours	(2) Off-Duty Hours
Number of Guilty	-0.208	-0.0691
Arrests (Last 3 Months)	(0.144)	(0.0465)
Number of Not Guilty	-0.233	-0.0245
Arrests (Last 3 Months)	(0.0652)	(0.0246)
Mean	18.883	18.890
Observations	45051	44927
Division FE	X	X
Year-Month FE	X	X
Badge FE		X

Standard errors in parentheses

Note: This table presents regressions of officer monthly off-duty hours worked on their number of recent arrests. The unit of observation is an officer-month. The arrest counts are moving sums from the previous three months. The first column includes fixed effects for officer division assignment and year-month. The second column additionally includes officer fixed effects. Standard errors are clustered at the officer level.

Figure A-1: Probability of Overtime Pay by Shift-Hour, Regression Adjusted

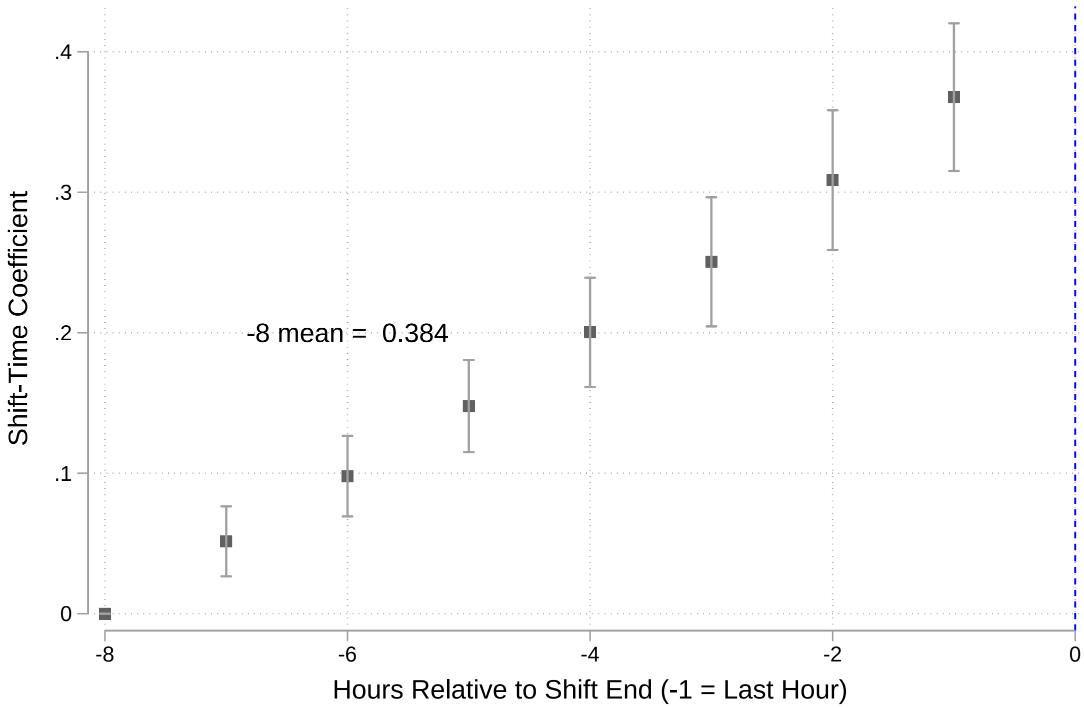

Note: Figure plots the probability that an arrest made in a given hour of an officer's shift leads to overtime pay, conditional upon the fixed effects in Equation (1). The -8 hour corresponds with the first hour of the shift; the -1 hour corresponds with the final hour of the officer's shift. 95 percent confidence intervals, clustered at the division-by-month level, are provided for each statistic.

Figure A-2: Histogram of Overtime Hours
Conditional on Last-Hour Arrest

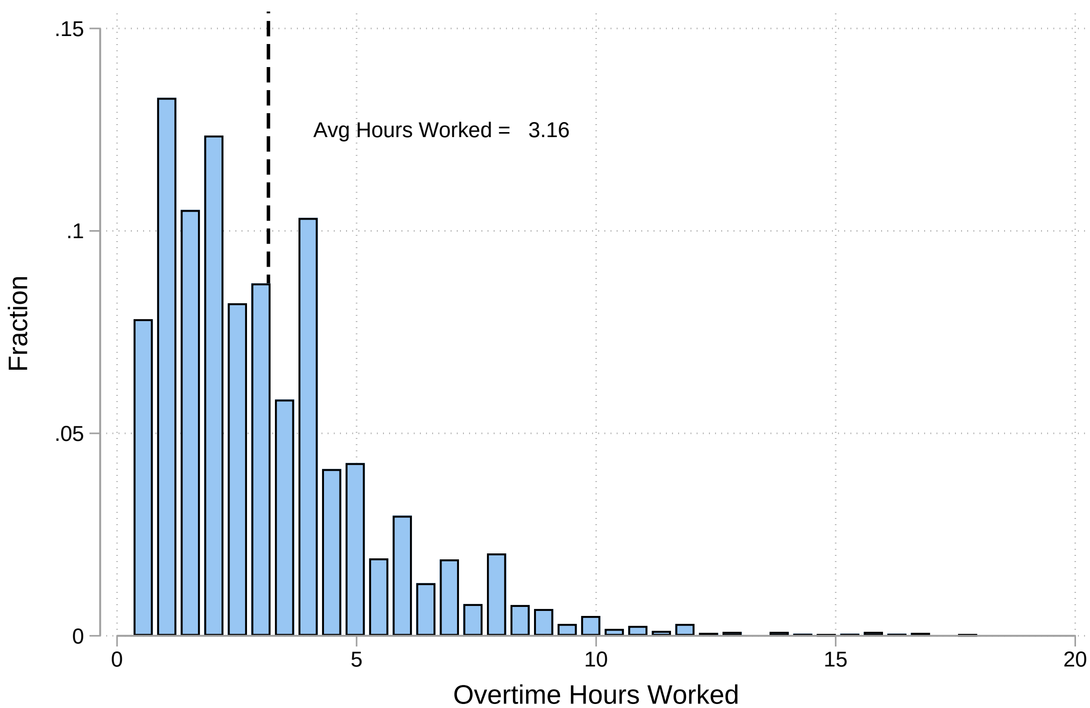

Note: This figure plots the histogram of hours of overtime payment received among overtime payments, conditional on officers making an arrest in the last hour of their shift and conditional on them receiving some overtime payment. The dotted vertical line plots the average overtime hours for this sample, which is 3.16.

Figure A-3: Arrest Frequency and Court Conviction Regressions

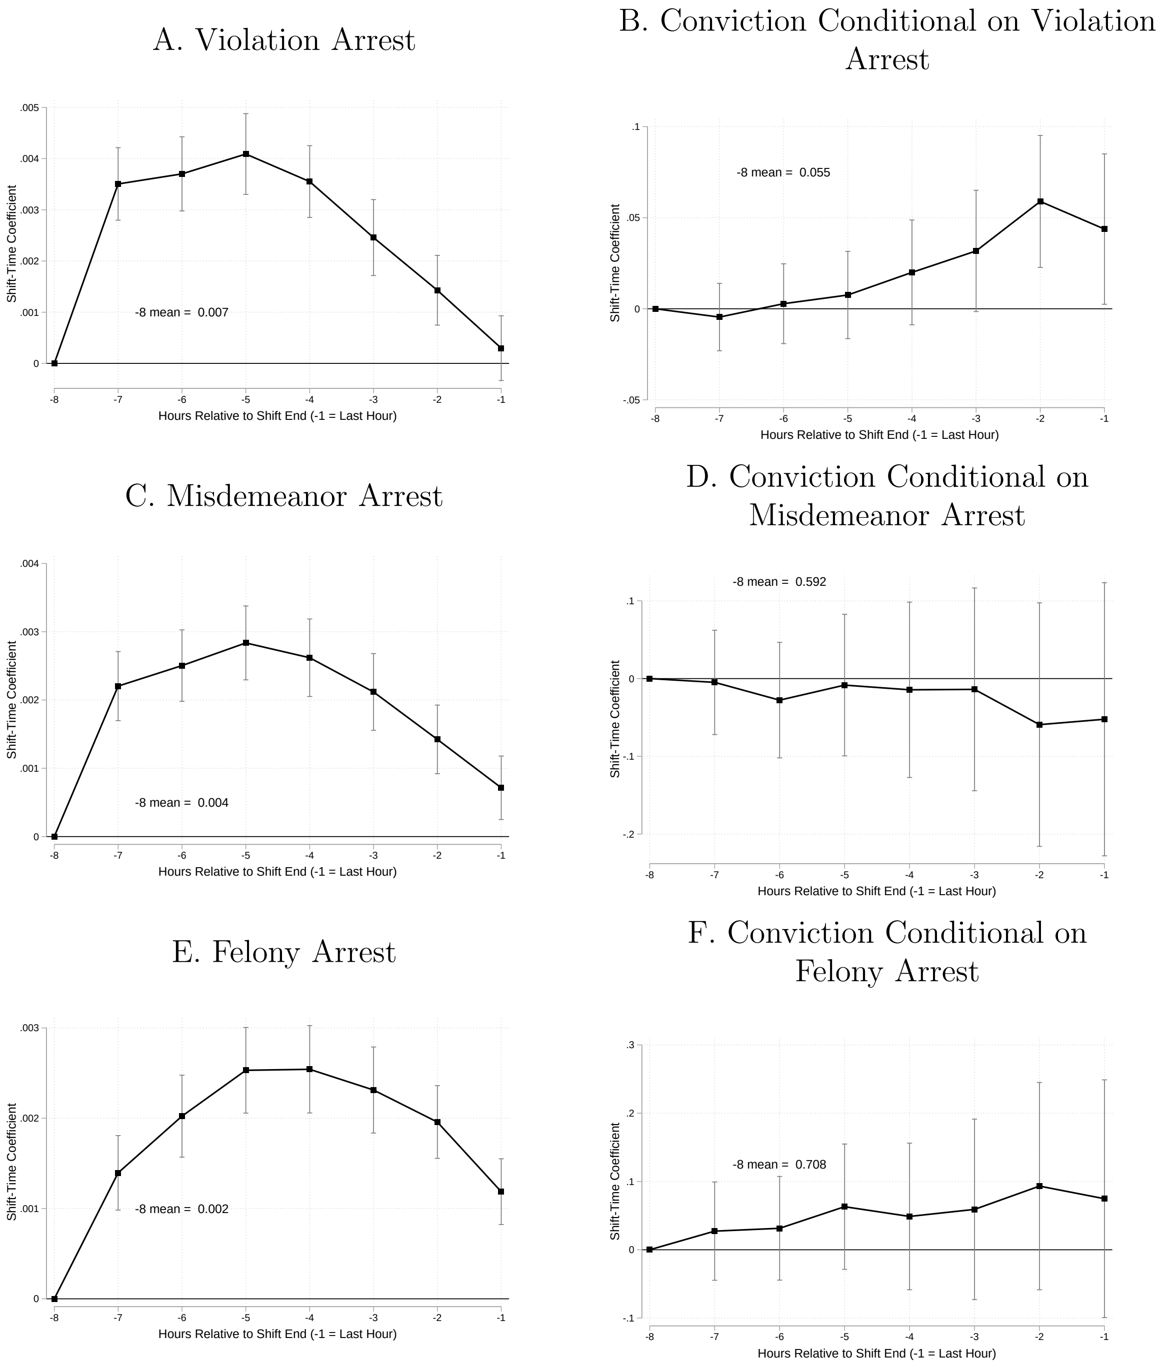

Notes: Figure plots each shift-hour coefficient from a regression of a given outcome variable on a vector of shift-hour indicator variables, conditional on the fixed effects described in Equation (1). We present estimates separately for arrests for violations, misdemeanor arrests, felony arrests and felony arrests excluding drug crimes. For each crime type, we also plot the probability of a conviction given an arrest. 95 percent confidence intervals, computed using standard errors that are clustered at the division-by-month level, provide a boundary around the point estimates. All coefficients are relative to the arrest incidence during the first hour of an officer's shift.

Figure A-4: Arrest Frequency and Court Conviction Regressions

A. Share Black Arrestee

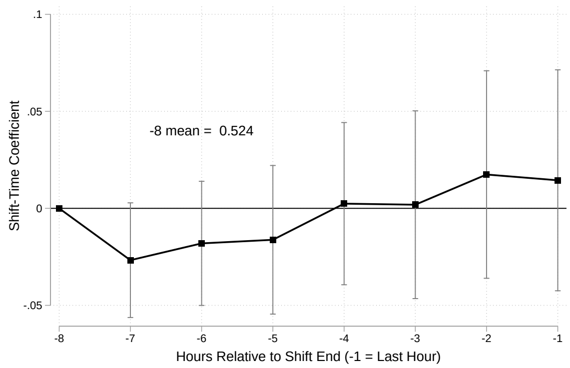

B. Conviction Conditional on Black Arrest

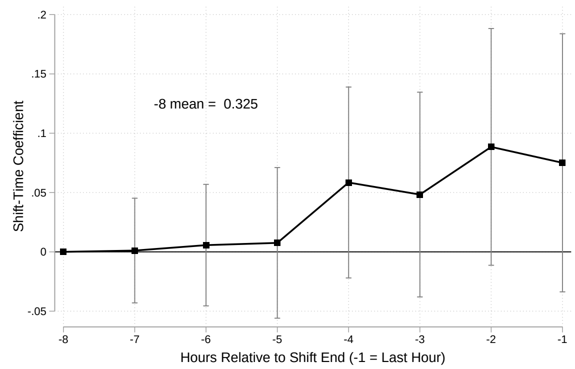

C. Share Hispanic Arrestee

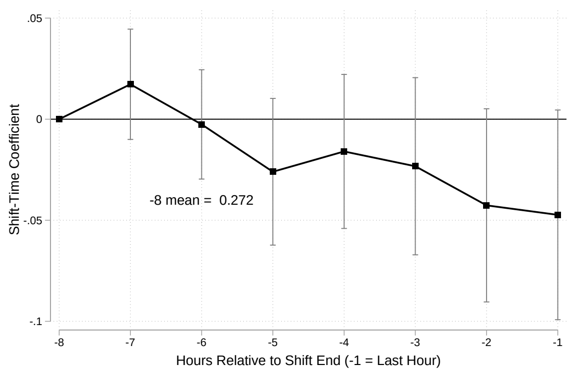

D. Conviction Conditional on Hispanic Arrestee

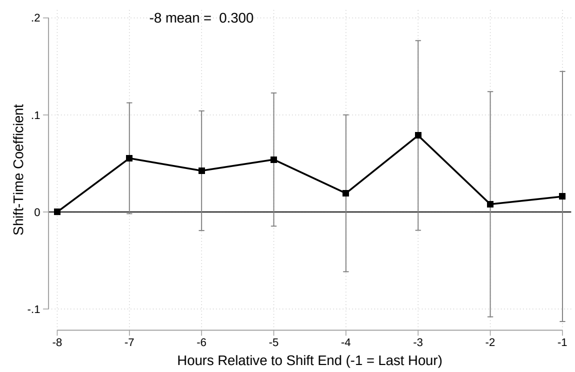

Note: Figure plots each shift-hour coefficient from a regression of a given outcome variable on a vector of shift-hour indicator variables, conditional on the fixed effects described in Equation (1). We present estimates for the share of arrestees who are Black and Hispanic as well as the conviction rate for Black and Hispanic arrestees. 95 percent confidence intervals are presented, computed using standard errors that are clustered at the division-by-month level. All coefficients are relative to the arrest incidence during the first hour of an officer's shift.

Figure A-5: Arrest Probabilities on Off-Duty Days

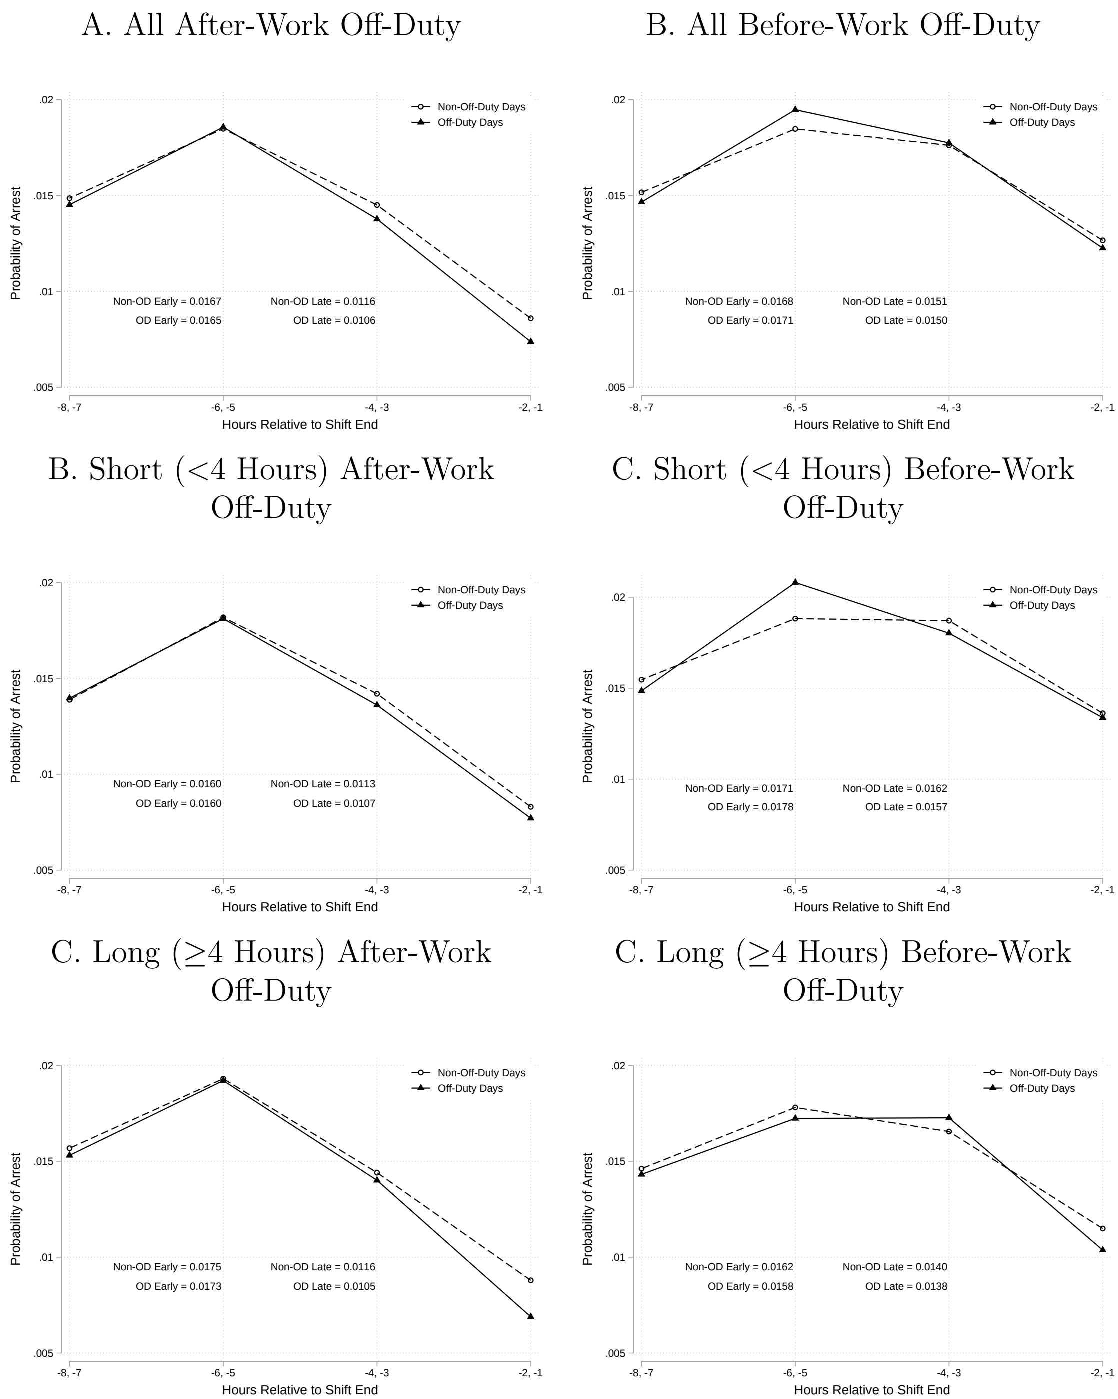

Note: This figure presents officers' arrest rates by shift hour, separately by the presence of an off-duty job. All panels group shift hours into two-hour time blocks. The left column of panels plots the arrest paths of officers who work a regular after-shift off-duty job, and the top left panel presents all such officers. In the bottom two left panels, we split this sample by whether the officers are working a short (<4 hours) or long (≥4 hours) shift. The right column of panels presents the analogous figures for officers who work a regular *before*-shift off-duty job.

Figure A-6: Visualization of Model

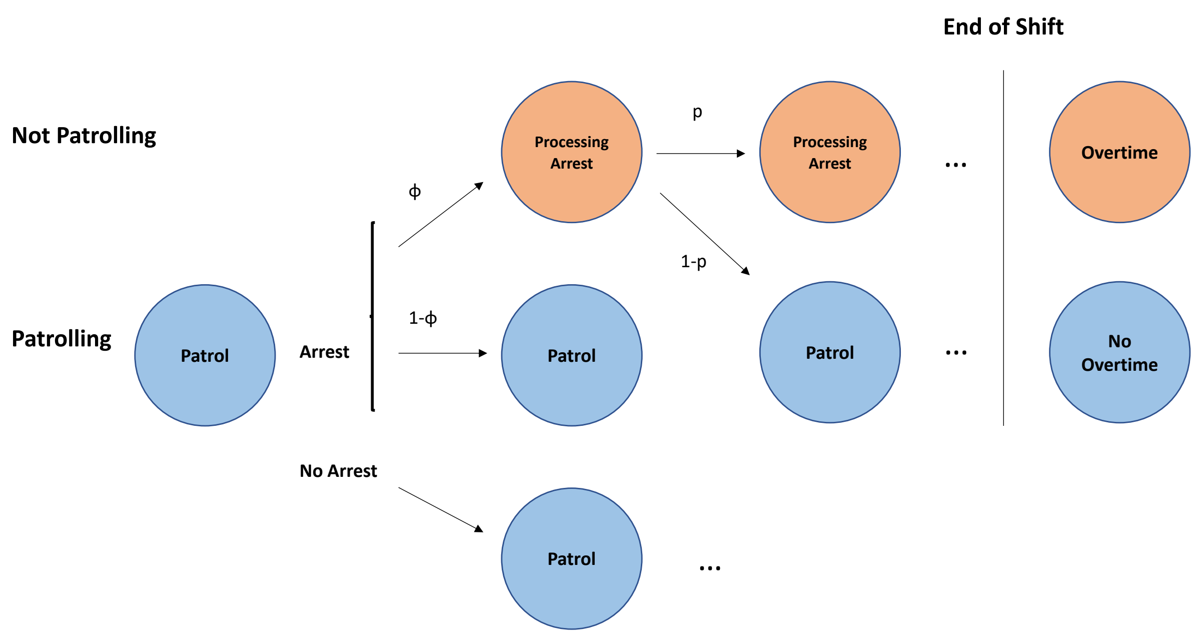

Note: Figure presents a schematic of the model presented in Section V.

Figure A-7: Model Fit, Off-Duty Guilt Effects

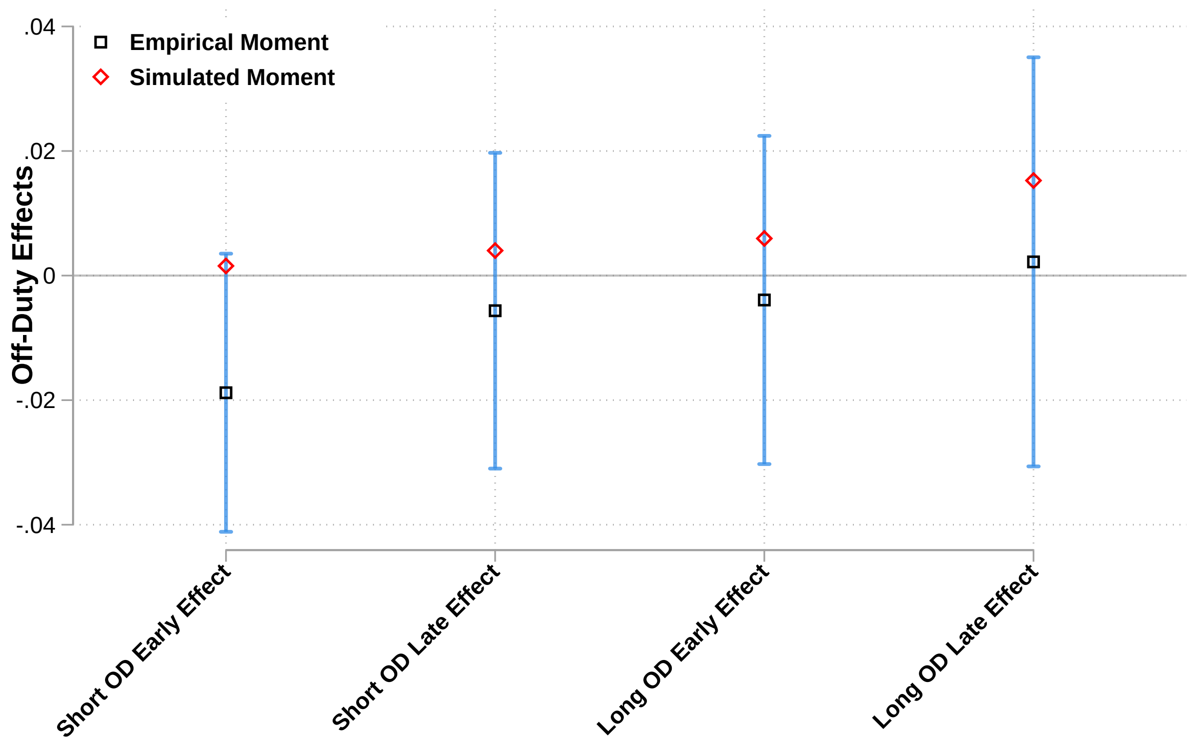

Note: Figure plots empirical versus simulated moments as a test of model fit for regular off-duty arrest impacts.

Figure A-8: Production Possibilities Frontier

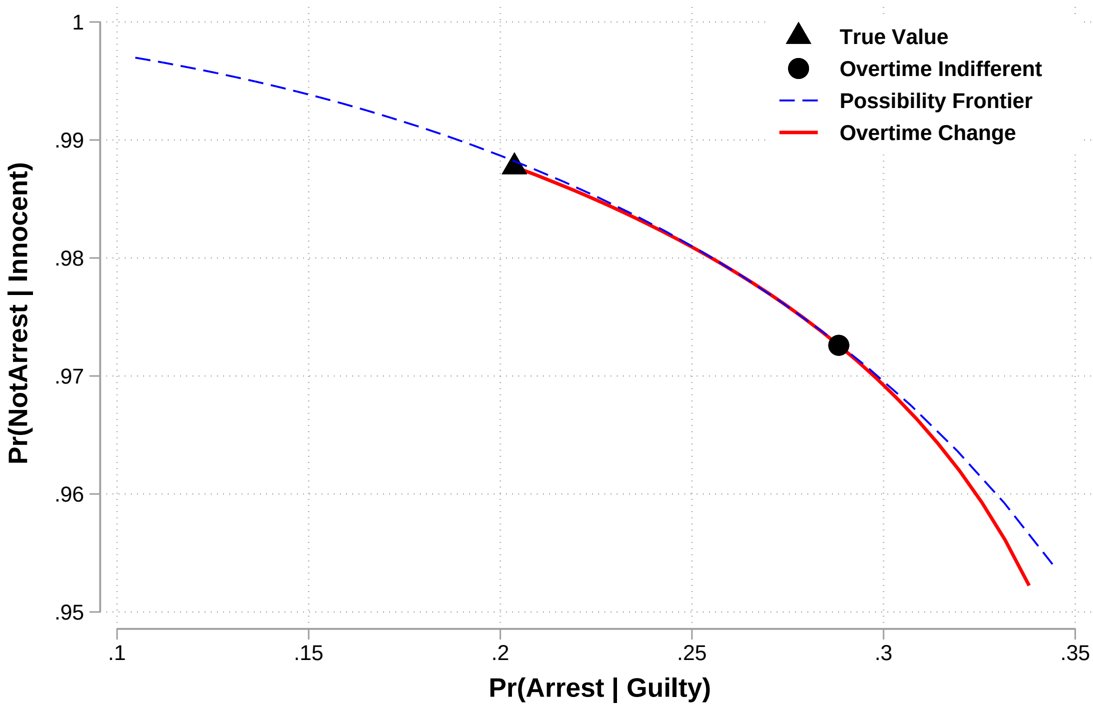

Note: This figure uses the model estimates to calculate the probability of not arresting a suspect who is innocent and the probability of arresting a suspect who is guilty for various parameter values, as described in Section V. The dotted line is constructed by setting $V_{ot} = 0$ (i.e. overtime indifference) and spanning different values of λ, the relative weights on arresting guilty suspects and not arresting innocent suspects. The solid triangle is the value for the baseline estimated parameters, and the solid circle is the value when V_{ot} is then set to 0. The red line traces the outcome for the baseline estimated parameters with progressive increases in overtime pay.

Figure A-9: Arrest and Court Regressions,
Officers Without Shift Changes

A. Arrests, Raw Plot

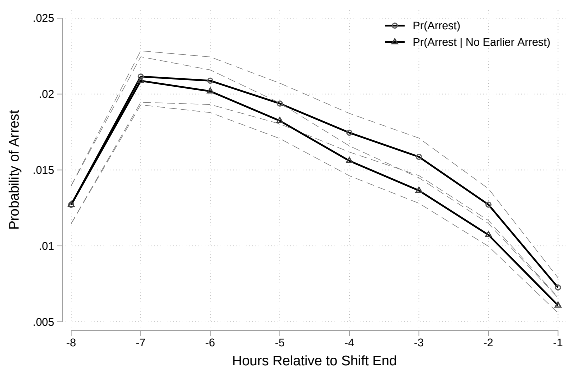

B. Arrests, Regression Plot

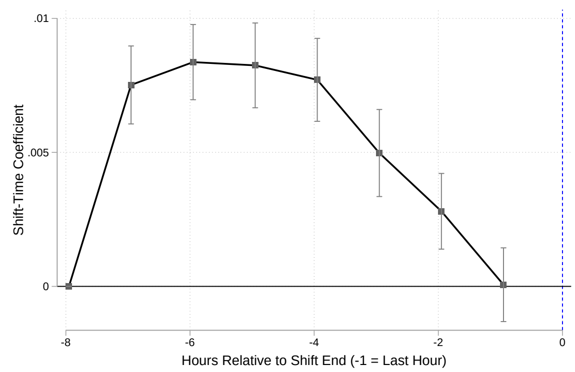

C. Court Conviction, Raw Plot

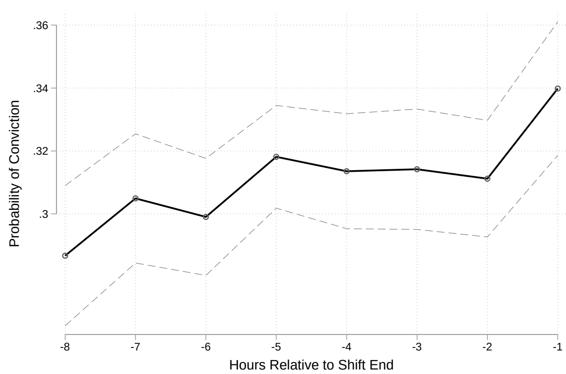

D. Court Conviction, Regression Plot

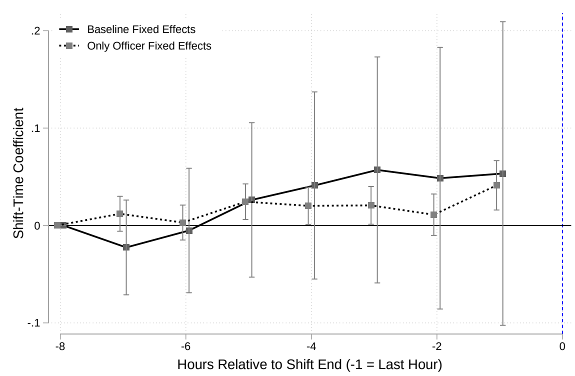

Notes: These figures reproduce the estimates from Figures 2 and 4, restricting attention to officers who do not change regular shift times or days during our analysis sample. In Panel D, we present regressions with both our baseline set of fixed effects and a specification with only officer fixed effects.

Figure A-10: Arrest and Court Regressions,
Officers With High Match Rate to 911 Calls

A. Arrests, Raw Plot

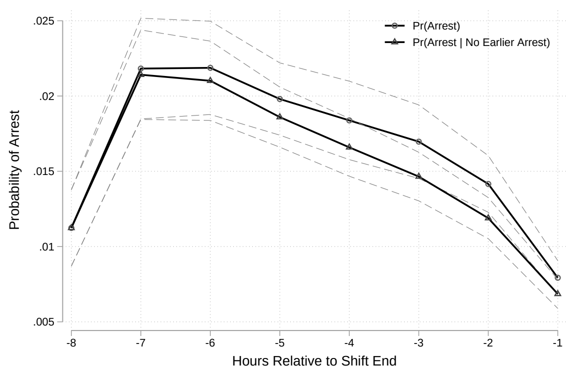

B. Arrests, Regression Plot

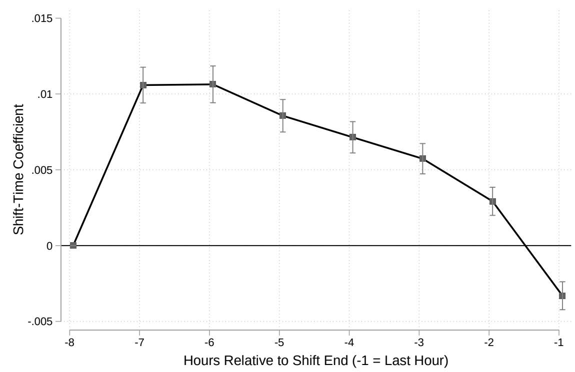

C. Court Conviction, Raw Plot

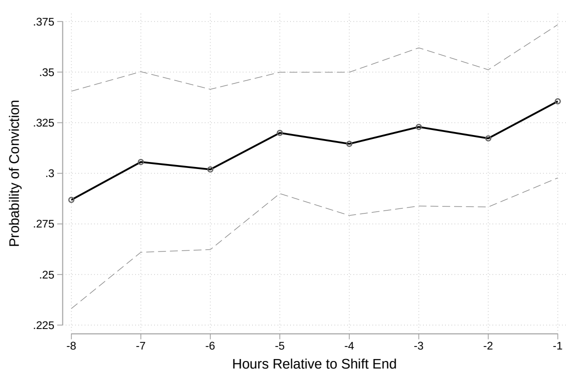

D. Court Conviction, Regression Plot

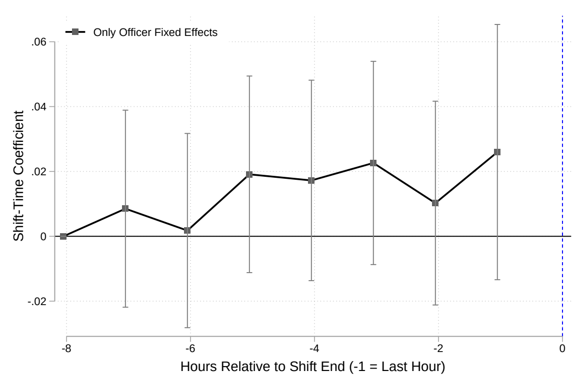

Notes: These figures reproduce the estimates from Figures 2 and 4, restricting attention to officers for whom at least 75% of their 911 calls match to their regular shifts.

Figure A-11: Arrest and Court Regressions,
Irregular Shift Days

A. Arrests, Raw Plot

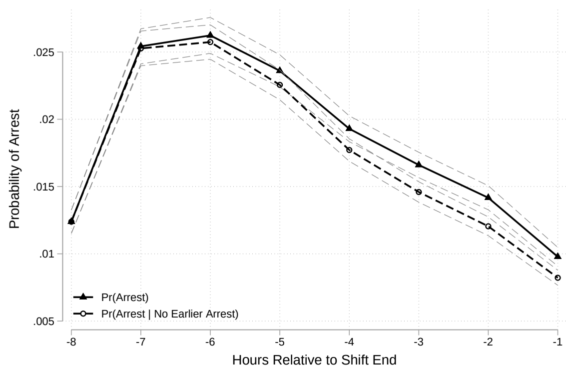

B. Arrests, Regression Plot

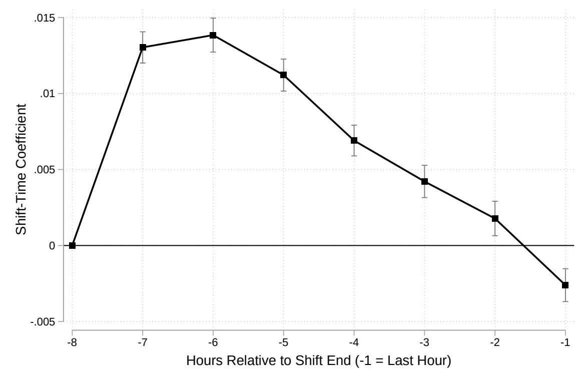

C. Court Conviction, Raw Plot

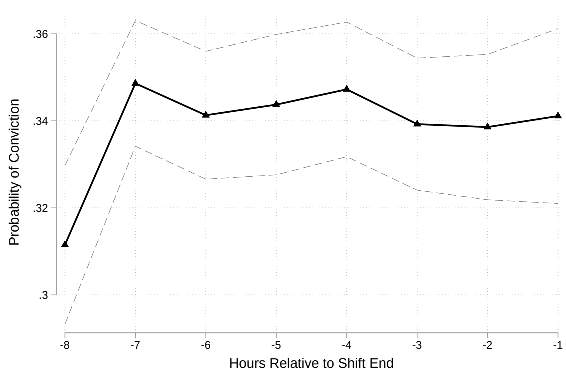

D. Court Conviction, Regression Plot

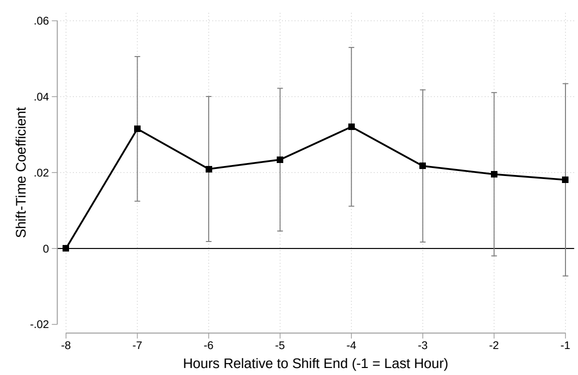

Notes: These figures reproduce the estimates from Figures 2 and 4, examining days when officers respond to 911 calls but are not working a regular shift. The construction of irregular shifts is described above.

Figure A-12: Arrest and Court Regressions,
Not Restricting to Shifts with 911 Call Taken

A. Arrests, Raw Plot

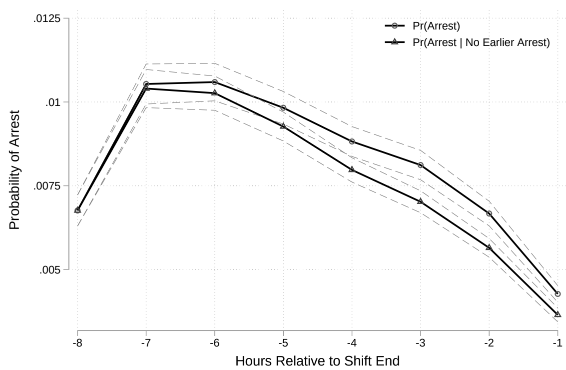

B. Arrests, Regression Plot

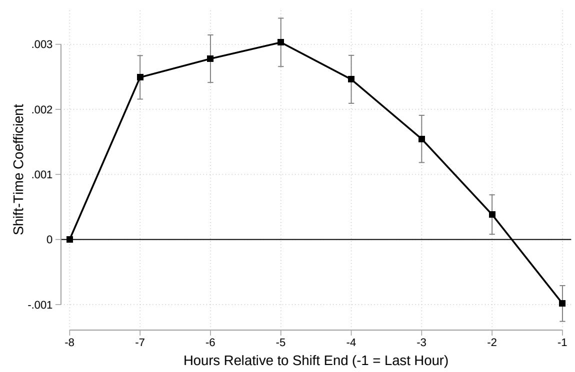

C. Court Conviction, Raw Plot

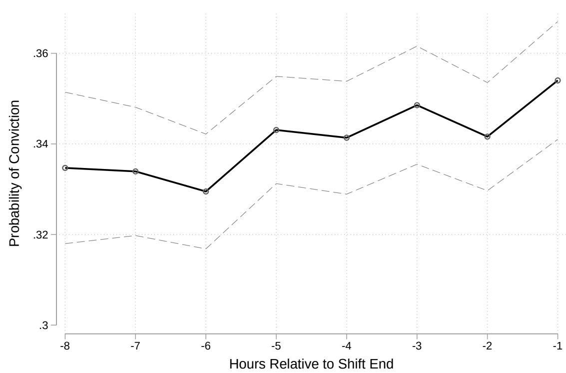

D. Court Conviction, Regression Plot

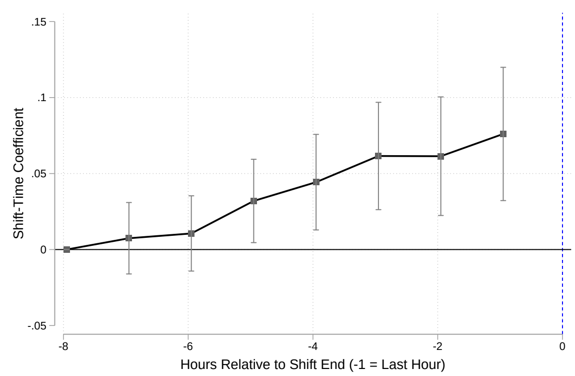

Notes: These figures reproduce the estimates from Figures 2 and 4, where we expand the sample to include shifts where an officer did not respond to a 911 call.

Figure A-13: Arrest and Court Plots,
Expanded Hours Around Shift

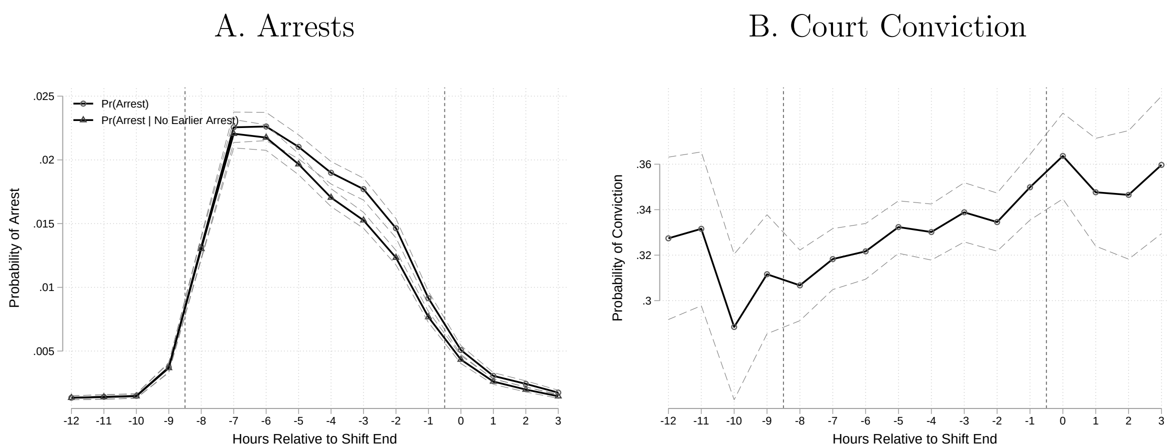

Notes: These figures reproduce the estimates from Figure 2, where we additionally include four hours before and after officers' shifts.

Figure A-14: Distribution of Officer Fixed Effects

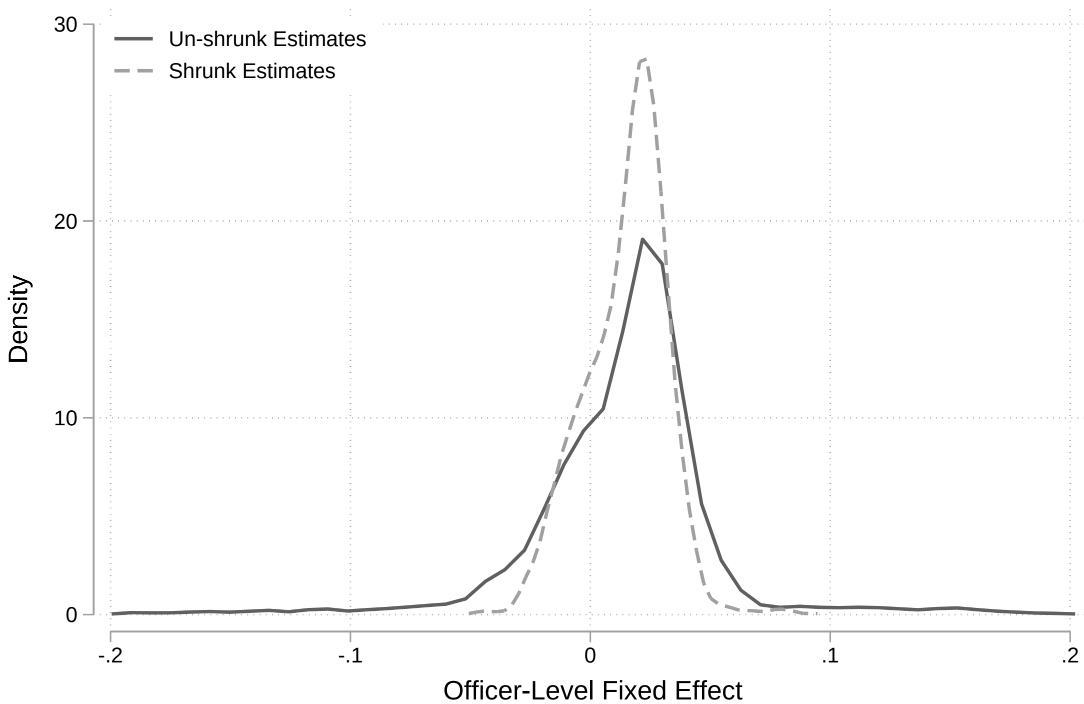

Note: Figure plots the the unadjusted distribution (the solid line) and the shrunken distribution (the dotted line) of officer fixed effects with respect to late-shift arrest activity. The fixed effects are estimated using a regression of whether an arrest occurs in the final two hours of an officer's shift on officer fixed effects. Because each officer makes a finite number of arrests, each fixed effect will be estimated with error and naturally some fixed effects will be estimated greater precision than others. To adjust the distribution for estimation error, we use a Bayes shrinkage approach similar to that employed in the teacher value added literature (Morris, 1983).

Figure A-15: Arrest Outcomes Across Officer Fixed Effect Percentile Cutoffs

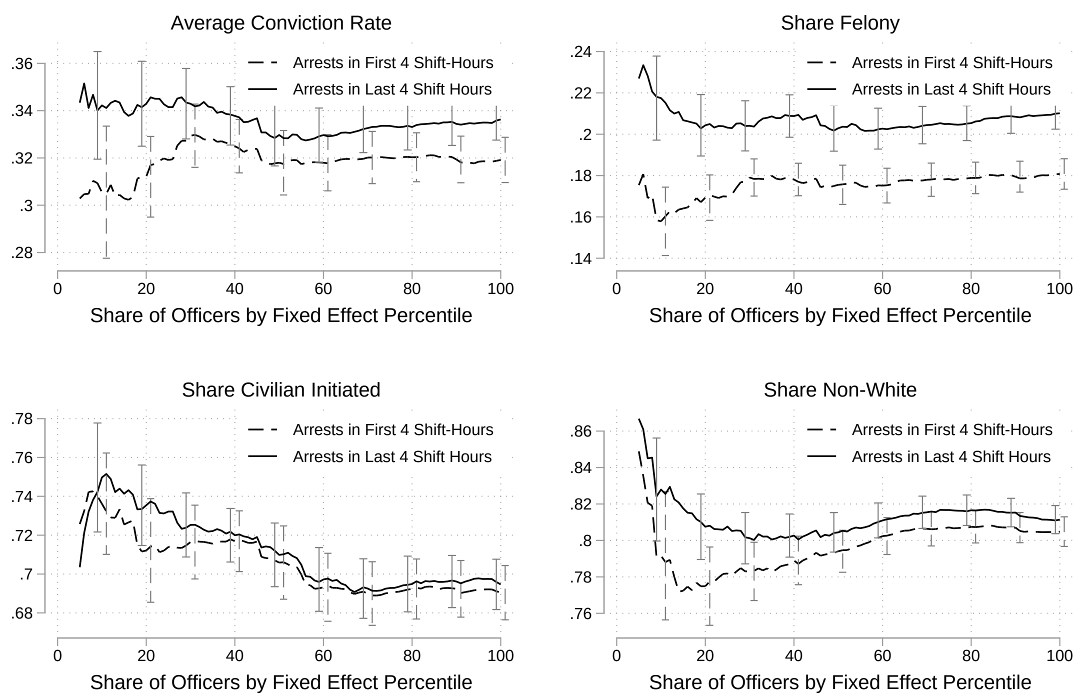

Note: Figure plots the mean of a given outcome variable separately for early shift arrests (the dashed line) and late-shift arrests (the solid line) for officers above a cutoff percentile of the distribution of the shrunken fixed effects. We present means for the conviction rate, the share of arrests that are officer-initiated, the share of arrests of non-white suspects and predicted overtime by crime type.

Figure A-16: Robustness to Alternative Specifications

A. Arrest Regression

B. Court Conviction Regression

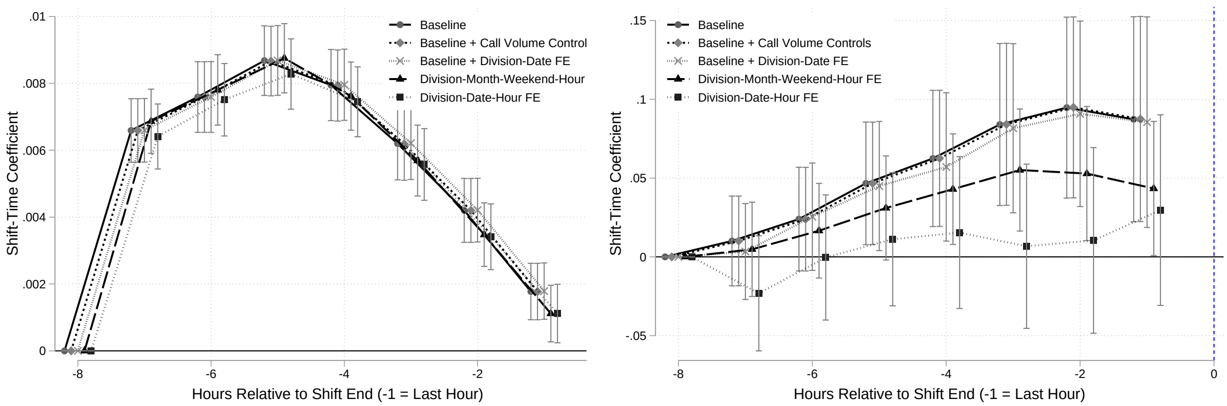

Note: Figures present versions of the Arrest and Court Conviction Regressions from Figure 4 with alternative sets of fixed effects.

Figure A-17: Accounting for Time-Into-Shift Effects

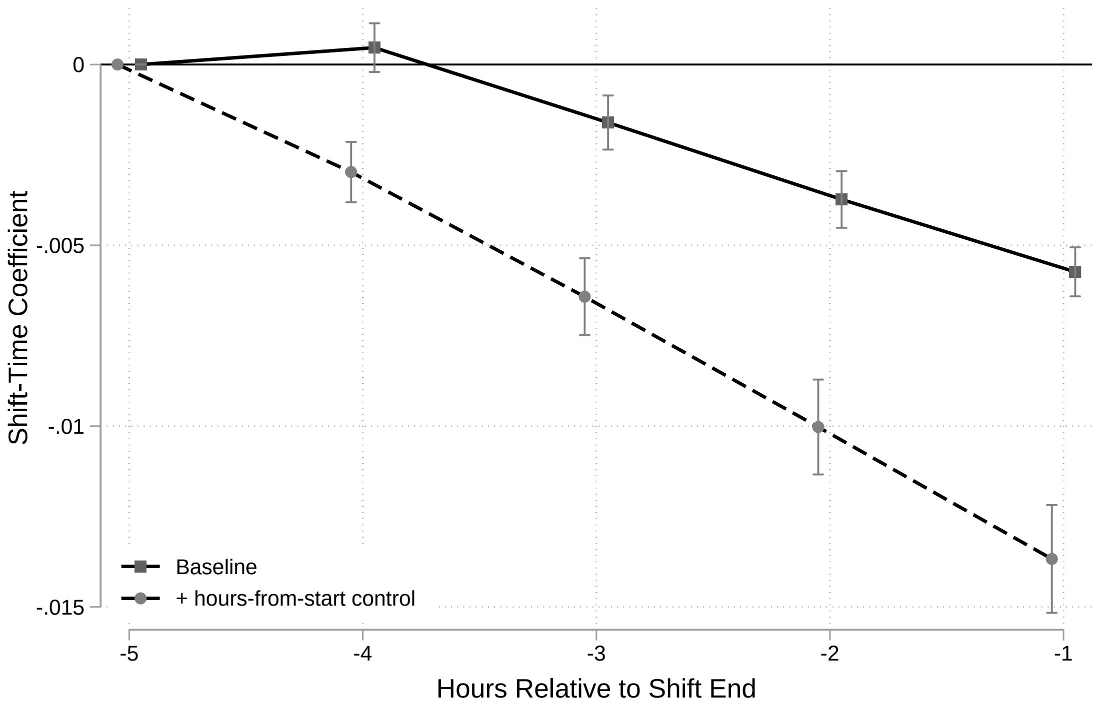

Note: Figure plots robustness analyses described in Section IV.C. The baseline sample is expanded to include nine-hour and ten-hour shifts. The "Baseline" coefficients are from estimation of Equation (1) with only coefficients for the last four hours, and the second set of coefficients are from the same equation with the inclusion of a variable for hours into the shift.

Figure A-18: Ruling Out Alternative Hypotheses

A. Number of Charges

B. Court Conviction Effect

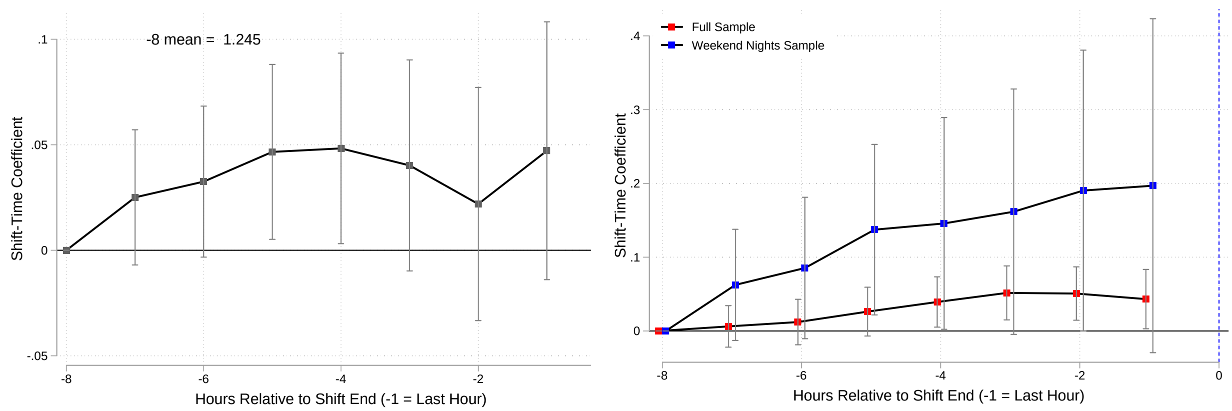

Note: The left-hand figure plots the number of arrest charges as a function of given hour of an officer's shift. The right-hand figure plots the probability of a court conviction in each hour of an officer's shift, separately for the full sample and for shift ending between 4:00pm and 10:00pm on Friday and Saturdays. The left-hand figure estimates condition upon the fixed effects in Equation (1). In the right-hand figure, our estimation includes fixed effects for division × hour × day-of-week, division × year × month, and badge, but we exclude division × shift × day-of-week because of reduced sample size. The -8 hour corresponds with the first hour of the shift; the -1 hour corresponds with the final hour of the officer's shift. 95 percent confidence intervals, clustered at the division-by-month level, are provided for each statistic.

Figure A-19: Model Counterfactual With and Without Fatigue Effect

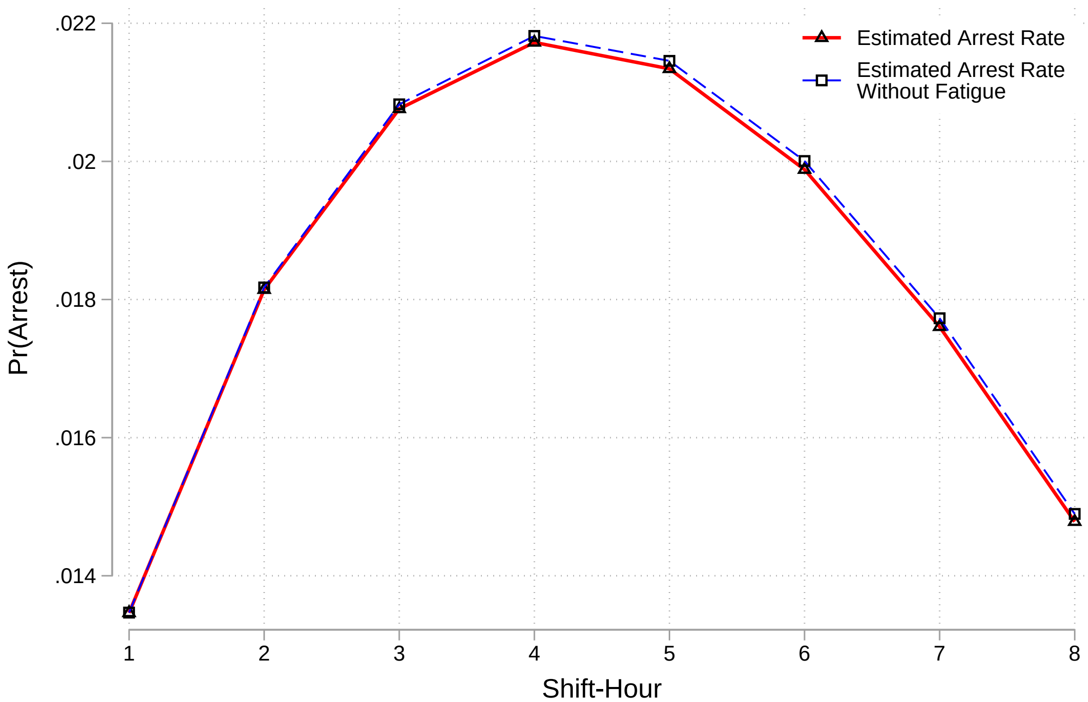

Note: This figure presents the model-estimated time path of arrest probability using the model specification from Section F that allows for a fatigue arrest cost. The solid red line presents the best fit to the observed data. The dotted blue line presents a counterfactual where the estimated fatigue cost is set to zero.

References

Linn, E. (2009). *Arrest Decisions: What Works for the Officer?* Number 5. Peter Lang.

Morris, C. N. (1983). Parametric empirical bayes inference: theory and applications. *Journal of the American Statistical Association 78*(381), 47–55.

Rehavi, M. M. and S. B. Starr (2014). Racial disparity in federal criminal sentences. *Journal of Political Economy 122*(6), 1320–1354.

Weisburst, E. K. (2022). whose help is on the way? the importance of individual police officers in law enforcement outcomes. *Journal of Human Resources*, 0720–11019R2.