

The Pro-Social Motivations of Police Officers*

Aaron Chalfin

University of Pennsylvania

Felipe Goncalves

University of California, Los Angeles

November 16, 2020

Abstract

How do public sector workers balance their pro-social motivations with private interests? In this study of police officers, we exploit two institutional features that change the implicit cost of making an arrest: arrests made near the end of an officer’s shift are more likely to require overtime work, and arrests made on days when an officer “moonlights” at an off-duty job after their shift have a higher opportunity cost. We document two consequences for officer behavior. First, contrary to popular wisdom, officers reduce arrests near the end of their shift, and the quality of arrests increases. We argue that these patterns are driven by officer preferences rather than departmental policy, fatigue, or incapacitation from earlier arrests. Second, officers further reduce late-shift arrests on days in which they moonlight after work, suggesting that they are, in fact, modestly responsive to financial incentives. Using these results, we estimate a dynamic model that identifies an officer’s implied trade-off between private and pro-social motivations. We find that police officers exhibit high pro-social motivation towards their work. In contrast to prior research showing that law enforcement outcomes are sensitive to financial incentives at an institutional level, the behavior of individual officers — the “street-level bureaucrats” who enforce the law — is not meaningfully distorted by monetary considerations.

JEL Codes: K42, J33, J45

*Chalfin: achalfin@sas.upenn.edu; Goncalves: fgoncalves@econ.ucla.edu. We are grateful to Bocar Ba, Amanda Geller, Sierra Kuzava, Ilyana Kuziemko, Adriana Lleras-Muney, Michael Makowsky, Conrad Miller, Emily Owens, Alessandra Voena, Till von Wachter, Emily Weisburst, and seminar participants at APPAM, SOLE, UCLA and the Arnold Ventures’ Police Accountability Roundtable for helpful comments. We thank the Dallas Police Department and the Dallas County Attorney General’s office for sharing their records and Lieutenant Stephen Bishopp and Corporal Julie Martin for help with institutional details. Any remaining errors are our own.

1 Introduction

A defining feature of public sector work is the scarcity of “high-powered incentives” in its employment contracts (Dixit, 1997, 2002; Banuri and Keefer, 2016). Instead, public sector employers often rely on hiring workers who are intrinsically-motivated towards their work and thus do not require strong financial incentives (Wilson, 1989; Heckman et al., 1996; Bénabou and Tirole, 2006; Tonin and Vlassopoulos, 2015; Besley and Ghatak, 2018; Linos, 2018). Government workers engage in activities that are multi-dimensional and difficult to monitor closely, and many bureaucracies thus face a particularly challenging principal-agent problem in which employees may take actions on behalf of the principal that maximize their own utility but are economically or socially counterproductive (Holmstrom and Milgrom, 1991; Finan et al., 2017). For these reasons, public sector performance is thought to both hinge on the employment of pro-socially motivated individuals and to be particularly sensitive to the distortionary impact of financial incentives when that process fails. This concern has been a principal factor in the design of public institutions in the developed world (Anechiarico and Jacobs, 1994; Glaeser and Goldin, 2006) and continues to hamper the provision of public services and basic governance in developing countries (Bowles and Garoupa, 1997; Mocan, 2008; Olken and Pande, 2012).

Despite the central importance of intrinsic motivation for the performance of public sector agencies, measuring the motivation of bureaucrats is a challenging task that is hampered by the availability of data and the difficulty of identifying natural experiments that can credibly elicit the true preferences of public employees. The majority of studies on the pro-social motivation of government workers use survey instruments to measure proxies of pro-social motivation, such as donations to charity (Tonin and Vlassopoulos, 2015) and working unpaid overtime hours (Gregg et al., 2011). While these are useful data points, it is unclear how closely these estimates correspond to the value agents place on their workplace activity. It is likewise difficult to use such estimates to understand the sensitivity of bureaucrats to the private incentives they face on the job and how such incentives might compromise the performance of their duties.

In this paper, we study the pro-social motivations of government agents in making a particularly high-stakes decision – the choice by police officers of whether to make an arrest. We assess the extent to which police officers place intrinsic value on their workplace activity relative to personal financial considerations and, crucially, whether changes to financial incentives can distort officer behavior. After a police officer makes an arrest, she typically must spend several hours processing the arrest, a requirement that can lead to working past the end of her shift and the receipt of overtime pay. While officers may place intrinsic value on arresting guilty suspects, their arrest decisions are also affected by the value they place on working overtime relative to their outside option. We evaluate how the quantity and quality of arrests vary throughout an officer’s shift and with the relative value of overtime work.

We do so by exploiting the fact that arrests made near shift-end have an increased chance of leading to overtime work and that arrests made on days where an officer “moonlights” in an off-duty job carry a higher opportunity cost. We then show how these estimates can be used to infer officers’ implied valuation for arresting guilty suspects and releasing innocent suspects (i.e., their pro-social motivations) and the value they place on working overtime relative to going home. More generally, our natural experiment allows us to identify the *corruptibility* of the typical police officer, which we define as the magnitude of the financial incentive needed for an officer to make an unambiguously incorrect arrest decision.

Because police productivity is difficult to measure and monitor, the potential for financial incentives to crowd out pro-social motivations is a particularly acute concern. Studies have found that police officers are more likely to issue traffic citations to out-of-state drivers when government fiscal conditions are tight (Makowsky and Stratmann, 2009, 2011) and that drug arrests tend to rise when asset forfeiture laws empower police to commandeer illegal assets (Benson et al., 1995; Blumenson and Nilsen, 1998; Baicker and Jacobson, 2007; Kelly and Kole, 2016).¹ Similarly, in its 2015 report on the Ferguson Police Department, the U.S. Department of Justice documented a systematic practice of excessive policing in minority neighborhoods which was carried out in order to raise revenue from fines for low-level violations (Department of Justice, 2015), a pattern of behavior which has sometimes been described as “policing for profit” (Blumenson and Nilsen, 1998; Skolnick, 2008; Worrall and Kovandzic, 2008). While the literature has documented numerous distortionary impacts of financial incentives at the institutional level, an equally important — and understudied — consideration is the private incentives of police officers, the “street-level bureaucrats” whose actions ultimately determine the enforcement of the law (Lipsky, 1971).

In recent years, a number of critics have raised the particular concern that overtime pay distorts officer activity towards making low-quality late-shift arrests (colloquially referred to as “collars for dollars”), a claim that has received attention in both the academic literature (Merola, 1981; Fyfe, 1984; Cavallaro, 1998; Moskos, 2008a; Punch, 2013) and popular media reports (Getter et al., 2014; Feuer and Goldstein, 2018; Olmstead, 2018). Despite growing public discourse around this issue, there has, to date, never been an empirical examination of the “collars for dollars” hypothesis. We address this question specifically and do so in a broader evaluation of how officers balance their intrinsic motivation with respect to workplace outcomes against their private preferences for overtime work.

We leverage a unique and unusually expansive set of administrative data from Dallas, Texas, which links information on officer shifts, overtime hours and off-duty work to 911 calls for service, arrests and court data. Critically, work shifts in the Dallas Police Department

¹In the related literature on corrections, researchers have found that the use of incarceration declines when counties bear the cost of incarcerating offenders (Ouss, 2015; Lofstrom and Raphael, 2016; Verma, 2016) and that for-profit private prisons house inmates for longer sentences than observationally similar inmates in public prisons due to the piece rate nature of compensation to private prisons (Mukherjee, 2017).

overlap throughout the course of the day, allowing us to distinguish between time-of-day and time-of-shift effects. We begin by documenting evidence that the probability of receiving overtime pay for an arrest rises monotonically throughout an officer’s assigned shift. When an arrest is made towards the end of a shift, officers are considerably more likely to collect overtime pay — in Dallas, 150 percent of the officer’s base wage rate — than when an arrest is made earlier in the shift or when no arrest is made at all. Given that an average arrest involving overtime wages requires at least three hours of an officer’s time, this financial benefit is not trivial; a late-shift arrest, on average, is worth an additional \$150 in salary or approximately two-thirds of an officer’s gross daily earnings.

We then make three central contributions. First, we document how enforcement activity changes throughout the course of a police officer’s shift. Despite the presence of strong financial incentives to make late-shift arrests, we find that the frequency of arrests *decreases* by approximately 30 percent towards the end of an officer’s shift. In addition, the quality of arrests, as measured through both the likelihood that an arrest leads to a court conviction and that it leads to a prison sentence, *increases* throughout the shift. We argue that these patterns are driven by officer preferences expressed through an increase in the threshold used to make an arrest rather than formal or informal department policies that discourage officers from taking overtime, the incapacitative effect of arrests made earlier in the officer’s shift, fatigue, and several other alternative explanations. While the decline in late-shift productivity suggests an aversion to working overtime (Chan, 2018), the increase in the quality of arrests made suggests that officers value arrest quality at the margin.

Second, we generate evidence on the degree to which police officers are responsive to *changes* in their financial incentives by documenting how officers respond to variation in the relative value of working overtime. In an ideal world we would be able to observe direct variation in the value of overtime pay. Because payment is fixed by the Fair Labor Standards Act at 150 percent of base hourly pay, we instead turn to an analysis of off-duty employment, work that is performed by police officers for private employers when they are not on-duty. Exploiting the fact that off-duty employment that is scheduled to occur just after an officer’s police shift increases the opportunity cost of overtime work, we find a significant though quantitatively small reduction in the propensity to make late-shift arrests on days when officers have a post-shift off-duty spell. In contrast, we find little evidence of a reduction in arrest activity on days where an officer works off-duty prior to their police shift. Taken together, these findings suggest that officers are, in fact, sensitive to financial incentives they face on the job, albeit to a modest degree.

Our third contribution is to feed these estimates into a simple dynamic model to estimate officers’ preferences over types of arrests, overtime work, and pay from overtime and off-duty work. We build on the theoretical work of Prendergast (2003, 2007) and model officers as having one margin of decision-making in each hour of their shift, which is the guilt threshold above which they arrest a suspect. Officers face a stochastic probability of needing

to process the arrest, which may both lead to overtime work and generate different choices of guilt thresholds throughout the shift. We identify the model by matching the regression-based time paths of arrest frequency, court conviction conditional on arrest, overtime pay for arrests in each hour of the day, and off-duty impacts. This final set of moments allows us to identify officers' dollar value for their arrest activity, since it shifts the opportunity cost of arrests by a specific dollar amount.

Our estimates suggest that officers derive a large intrinsic value from making accurate arrest decisions. Specifically, they place a value of \$1,382 on correctly not arresting an innocent individual and a value of \$7,621 on correctly arresting a guilty individual. These values of pro-social motivation are, to our knowledge, the first dollar-denominated estimates in the literature. Critically, these values are large relative to the magnitude of income derived from overtime pay and off-duty work, suggesting that police behavior is unlikely to be significantly distorted by the presence of either of these two incentives. Indeed, our estimates also allow us to consider the potential *corruptibility* of officers, which we define as the magnitude of financial incentive needed to make either a clearly wrongful arrest or fail to make a clearly justified arrest. An officer in the last hour of her shift would need to expect at least \$3,548 in overtime pay to make an arrest of an individual she knows with certainty is innocent. Conversely, an officer in her last hour of the shift would need to have at least \$8,024 of income waiting at an off-duty job to forgo arresting an individual she knows with certainty is guilty.

We then use the model estimates to conduct a number of simple counterfactual exercises. Consistent with our corruptibility estimates, even large increases in overtime pay lead to only modest changes in behavior — the elasticity of arrests with respect to expected overtime pay is approximately 0.03. Perhaps more interestingly, when we allow officers to be indifferent between working late and going home (a purely “altruistic” officer), arrest propensity increases throughout the shift and arrest quality declines. These effects are driven by the fact that early arrests carry the opportunity cost of precluding officers from making later arrests. This striking result — that altruistic officers appear to be engaging in “collars for dollars” — highlights the importance of formally modeling the dynamic decision that officers face in order to infer their preferences.

Several features of our setting are ideal for measuring pro-social motivation and the preferences of public sector agents more broadly. First, our data are exceptionally detailed, allowing us to observe individual behavior at an hourly level. Measures of workplace behavior are typically based on subjective assessments (such as promotions or performance evaluations), where systematic reporting bias is a threat to identification. By assessing changes to behavior that are not dependent on performance evaluations, we avoid this concern. Second, we use a shift in the cost to a workplace activity that can be denominated in dollar terms. Because we can provide a plausible estimate for the increased opportunity cost of making an arrest when an officer has scheduled off-duty work, we can estimate the dollar

value placed on making an arrest of a guilty individual and avoiding an arrest of an innocent individual. Third, due to historical concerns about police corruption (Walker, 1977; Greene, 2000; Potter, 2013), the public sector agents we study have relatively low-powered incentives related to their workplace activities. As in most U.S. cities, promotions in the Dallas Police Department are exam-based, where officers are ranked on their performance and promoted down the list as positions appear. This process does not incorporate information on arrest activity. Similarly, the base pay for officers is a function of their rank and experience and is not tied to their arrest activity beyond the impact on overtime pay. Because of these institutional features, officers do not receive professional benefits from their arrest activity, making our setting uniquely suited to credibly identifying pro-social motivations.

Our paper contributes to the large literature on the personnel economics of public sector agencies (Prendergast, 2001, 2007). Several papers have documented the impact of changes to workplace incentives on the performance of public sector agents. For example, Mas (2006) finds that, after losing a collective bargaining arbitration case, police officers significantly reduce their enforcement activity. Similarly, Bertrand et al. (2020) find that, when future promotion prospects are worsened, bureaucrats reduce their performance. Most close to our study, Ash and MacLeod (2015) show that, when state court judges have their work obligations reduced, their written opinions are better cited in the future, suggesting an increase in work performance and an intrinsic motivation for performing their job well. In our context, we show that the work performance of police officers is influenced by their private cost to making an arrest. We advance the literature by providing the first dollar-denominated estimates of intrinsic motivation, identified through changes in workplace incentives. These estimates allow us to answer a wide range of questions about the distortionary impact of monetary incentives.

This study also contributes to the literature on the labor supply of workers with shift-length discretion (Oettinger, 1999; Farber, 2015). In particular, our paper is closely connected to Chan (2018), who finds that emergency department shift schedules induce physicians to “slack off” at the end of their workday by accepting fewer patients near end of shift and spending less time with patients that they do meet with. Similar to physicians, we find strong evidence that police officers place great value on their leisure time. In our context, given the intrinsic value that officers also place on arresting guilty suspects and avoiding false arrests, incentives created by overtime pay are insufficiently large to change police decision-making at the margin.

The remainder of the paper proceeds as follows: Section 2 provides institutional background on police patrolling, arrests, and overtime. Section 3 provides a description of the administrative data from the Dallas Police Department and the court data and how the data were linked together, Section 4 describes the empirical models, Section 5 discusses the results, Section 6 presents the dynamic model and estimates of officer preferences, and Section 7 concludes.

2 Institutional Background

While their actions are constrained by both law and departmental policy, police officers typically have broad discretion over whether or not to make arrests (Goldstein, 1963; Kelling, 1999; Mastrofski, 2004; Groeneveld, 2005; Linn, 2009; Bronstein, 2014; Owens, 2020). Among patrol officers, discretion stems primarily from two sources. First, officers must inevitably use their judgment to determine whether they have probable cause to make an arrest or whether they have reasonable suspicion to search a suspect, a condition which is sometimes an antecedent to an arrest (Mastrofski, 2004; Goel et al., 2016). This source of discretion means that arrest decisions can depend on an officer’s subjective perception of guilt and their understanding of legal concepts about which there is often substantial ongoing jurisprudence. Second, police officers can decide how proactive they would like to be in searching for and identifying criminal activity (Sun, 2003; Wu and Lum, 2017). Proactive policing is particularly relevant for crimes like drug violations or weapons possession, for which there is no victim and which therefore must be discovered by law enforcement. This type of discretion likewise creates numerous opportunities for officers to engage in selective enforcement in order to potentially take advantage of the availability of overtime pay.

Processing an arrest typically requires significant police effort. In all cases, officers must prepare a detailed report describing the arrest and the associated incident. Typically a suspect must be taken into custody, which means that the officer has to transport the individual to a station to be booked and fingerprinted.² The jail intake process is lengthy, and officers are required to remain with the suspect during the entire booking process. Depending on the charge, some arrests require additional steps, such as blood alcohol testing for drunk drivers.³ Likewise, processing time may be further extended if the arrestee requires medical care. In sum, even arrests for low-level charges typically require a minimum of several hours of work.⁴

If processing an arrest requires that an officer must work past his or her regular shift time, the officer will receive overtime compensation at a rate of 150 percent of their base hourly pay rate, as stipulated under the Fair Labor Standards Act. Officers can choose whether to

²While arrests are generally associated with being jailed, arrests for low-level infractions (i.e. misdemeanors or violations) sometimes are issued in the form of a citation with a court date. In these cases, an officer does not need to book the individual and will likely require less time to complete processing.

³Because we focus on city patrol officers, traffic infractions are a small share of cases – just 3 percent of cases are for driving while intoxicated.

⁴After the initial period of processing, arrests may lead to court time, since charges that lead to a criminal case typically require the officer to provide testimony. While the vast majority of arrests either do not result in a conviction or result in a plea agreement soon after an arrest, court overtime potentially remains a motivating factor in making arrests. Our identification strategy nets out this motivation, focusing solely on changes in the incentive to arrest late versus early in an officer’s shift.

receive direct pay or compensatory time, in which case they receive time and a half that can be used at their discretion. While in some cities, arrests may not lead to overtime if certain steps in processing are not immediately available — for example, if officers must send their paperwork to the district attorney’s intake office, which is closed at night — in Dallas, all steps typically must occur immediately after the arrest has been made.

Do officers seek to secure access to overtime pay by making late-shift arrests? We might expect arrests to be more common at precisely the time when the expected overtime payment for an arrest is highest. However, because overtime arrests are not costless for an officer, such a story need not be true. Indeed, if police officers, like other workers, value both income and leisure, it is *a priori* unclear whether an officer will prefer to make an arrest at the end of a shift in order to secure overtime pay or if he would otherwise prefer an earlier shift arrest in order to maximize leisure time. As noted by Moskos (2008b) and Linn (2009) among others, a variety of anecdotal evidence supports the view that private motivations may often be inconsistent with a desire to maximize overtime pay. In particular, officers may have a variety of personal reasons to avoid making late-shift arrests including a “hot date,” a sick baby, a college class or, critically, an off-duty job, all of which would be considerably disrupted by a late-shift arrest (Linn, 2009). Likewise, police officers, like other public sector workers, may be highly intrinsically motivated and therefore find it costly to make an unjust or socially undesirable arrest. The extent to which officers make late-shift arrests will ultimately depend on the relative value they place on their workplace activity (i.e., their pro-social motivation), overtime pay, the ardour of overtime work, and their planned activities after their shift.⁵

3 Data

Data were obtained from a series of public records requests to criminal justice agencies in Dallas, Texas and cover information for the years 2015 through 2019. Arrest records come from the Dallas Police Department (DPD). These data provide information on all arrests made by DPD officers and include the date, time, location, and all of the arrest charges. The arrestee is identified by his or her full name and age at the date of the offense, and the arresting officers are identified by their badge numbers.

To identify whether an arrest leads to a criminal conviction or a sentence, we use criminal court data from the Dallas County Attorney General’s office covering 2015 to 2017. Each case reports defendant’s full name, date of offense, criminal charges, and final case disposition. We link these records to our arrest data using a fuzzy match on first and last name and

⁵Conversations with personnel at the Dallas Police Department reveal no indication that there are either formal or informal limitations on the amount of overtime work that officers are allowed to perform. This is consistent with the available data which suggest that there is a great deal of heterogeneity among police officers with respect to the number of overtime spells worked.

offense date (Lahiri and Larsen, 2005; Tahamont et al., 2020).⁶ We consider an individual to be guilty if their case is not dismissed and they are not found innocent by judge or jury. We likewise include as guilty cases that end in a plea of “non-adjudication of guilt,” where an individual does not contest the charge but does not formally admit to guilt. Our measure of a criminal sentence corresponds to whether any prison or jail sentence is reported in a case, including cases where the individual is assigned a sentence probation in lieu of serving prison or jail time. In cases in which the arrest does not link to any court records, we conclude that the suspect was not found guilty and was therefore not sentenced.⁷ Our main measures of arrest quality will be court conviction and sentencing, an approach we borrow from previous studies in criminology (Brewer and Gilliam, 1979; Forst, 1982) and economics (Ater et al., 2014; Weisburst, 2020). While convictions and sentencing may not perfectly capture guilt, Cassell (2018) and Loeffler et al. (2019) find that wrongful imprisonment, while highly costly, appears to be surprisingly rare — occurring in an estimated six percent of cases.

To link every arrest to the relative time during an officer’s shift when the arrest was made, we use a record of officer overtime payments from the police department. These data record the date, number of hours, and total payment of the overtime spell. Each entry also include an officer’s regular assignment location, regular shift hours, and days of week on which they work, which we use to calculate the number and range of hours in which an officer is regularly working. In our sample period, 5,327 officers have made an arrest, and 49.9 percent of these officers appear at least once in the overtime data.

To identify arrests that were precipitated by a citizen call for service, we use data on 911 calls for service from the Dallas Police Department. For each call, these data record the date, time, location, call description, responding officers, and incident numbers for any reports written from the call. Through the incident number, we link these data to arrest records to identify whether an arrest is made. We use these data to identify officers who are working patrol on a given day and are used in our robustness section to restrict attention to arrest propensity from calls.

To measure secondary work activity, we use a database of all days when an officer has an

⁶We link all arrests and court records by first and last initial, remove links where the first and last names deviate by more than two characters or the offense dates are more than two days apart, and then keep the strongest link for each arrest record. We are fairly confident about the strength of the match. For the set of all first and last initial matches, there is a large spike of cases where the offense dates coincide. When the offense dates and initials match, the majority of names have no disagreement between the two data sources. So our data matching aligns fairly closely with an exact matching algorithm.

⁷The court records comprise all cases seen in the county criminal court system. This system oversees all cases with charges of a Class C Misdemeanor or higher. However, it is not uncommon for arrests with a lower charge in our data to appear in the county court data. Among Violations (below misdemeanor), 10 percent of arrests are linked to a case in the county court. We therefore construct the indicators for guilt and sentenced for all of our arrests and note that the outcome measures guilt but also to some extent reflects the severity of the arrest.

off-duty job for 2015 to 2018. These records include the employer, date, shift start and end time, and officer name and badge. Henceforth, we refer to these assignments interchangeably as off-duty, moonlighting, or secondary jobs.

3.1 Sample Construction

The unit of observation for our baseline sample is a date and hour in which an officer is working his or her regular shift. From the overtime data, we observe the regular shift hours and days of work for each DPD officer. We restrict attention to officers whose listed assignment is one of the seven patrol divisions or the central business district, excluding officers working in the traffic unit and other specialized units such as narcotics, violent crimes, and tactical response and support. We further exclude all days-hours where an officer's two closest overtime payments list different regular shift assignments, hours, or days off. To avoid days where an officer may have called in sick or is otherwise not working patrol, we keep only shifts where an officer appears in the 911 data as taking at least one call. We also restrict attention to days where the officer's shift is exactly eight hours in length, removing the small share of days where an officer's shift is nine or ten hours. We do so because these days have a significantly lower arrest rate, indicating that officers are conducting qualitatively different activity on these days.

3.2 Summary Statistics

In **Table 1**, we present descriptive data on the 6,094 DPD officer-years in our data. The average DPD officer has served approximately 12 years on the job. The department is fairly diverse — 26 percent of officers are black and 20 percent are Hispanic. The average officer in our sample earns just over \$64,000 per year in base salary and earns an additional \$5,717 in overtime pay. While overtime pay accounts for approximately 9 percent of an average officer's compensation, receipt of overtime wages varies considerably among officers ranging from zero to more than \$127,000 (SD = \$9,004). Given the broad range of overtime pay among officers, the data are consistent with our understanding of policy and practice at the Dallas Police Department, which does not have a formal cap on overtime pay. To the extent that informal department policies might discourage officers from working overtime spells, these constraints appear to be mostly non-binding.

On average, patrol officers respond to 620 emergency calls per year (minimum = 1, maximum = 2,511). Officers make, on average 27.8 arrests per year, 37 percent of which resulted in a misdemeanor or a felony conviction. 64 percent of arrests by the average officer are initiated by a civilian call for service; the remainder are initiated by the police officer.⁸ **Table 2** provides descriptive statistics on the types of arrests that DPD officers

⁸We identify arrest origin through an incident report or 911 call, which can be linked to the arrest data through incident number and name the complainant. When the incident is officer

make. With the exception of warrant arrests, the most common type of arrest was for drunk and disorderly conduct (15 percent), followed by assault (13 percent), narcotics violations (11 percent), retail theft (6 percent) and public intoxication (5 percent). Arrests for violent offenses are rare — just 3 percent of arrests are for either felony or misdemeanor assault and under one percent are for robbery, rape or homicide. Overall, 18.1 percent of arrests were for felony offenses. As one might expect, the conviction rate varies considerably by the type of arrest. Some arrests are incredibly likely to lead to a conviction — conviction rates are extremely high for driving while intoxicated (83 percent), trespass (83 percent), narcotics violations (80 percent) and retail theft (75 percent). On the other hand, convictions are rare for disorderly conduct (7 percent), the most common arrest type in our data. An average arrest has a 45 percent chance of leading to an overtime spell. This figure is slightly higher (54 percent) for DWI arrests, the fourth most common arrest type in the data and is somewhat lower than average for low-level offenses such as public intoxication and disorderly conduct.⁹

3.3 Descriptive Analysis

Prior to introducing our empirical strategy, we present descriptive evidence from the raw data. We begin by establishing that the probability of receiving overtime pay varies by the shift-hour when an arrest is made. In **Figure 1**, we plot the probability of receiving overtime pay (Panel a) and the number of overtime hours conditional on overtime receipt (Panel b) against the shift-hour of arrest. The -8 hour corresponds with the first hour of an officer’s shift; the -1 hour corresponds with the final hour of the officer’s shift. In each graph, the two black lines plot the probability of receiving overtime pay for an arrest made in each shift-hour — one of the lines corresponds with felony arrests; the other for non-felony arrests which are comprised of both misdemeanor arrests or violations.

In Panel (a), the blue line indicates that the probability of receiving overtime pay when no arrests are made is approximately 29 percent which reflects the fact that recent arrests are not the only reason why officers work overtime hours. Nevertheless, the probability of working an overtime spell is considerably higher on days in which an officer makes an arrest. The probability of receiving overtime pay having made an early-shift arrest is approximately 36 percent for a non-felony arrest and 45 percent for a felony arrest, and there is a near monotonic increase in the probability of overtime receipt throughout an officer’s shift. During the final three hours of an officer’s shift, the probability of receiving overtime pay is approximately 75 percent for a felony arrest and approximately 65 percent for a non-felony arrest.¹⁰

initiated, the complainant is usually listed as “City of Dallas.”

⁹This is to be expected as processing a DWI offender typically will involve waiting for a warrant to draw blood as well as the blood draw itself.

¹⁰In **Appendix Figure 2**, we present the analysis having conditioned on officer, division \times

Panel (b) plots the expected number of overtime hours received conditional on working any overtime, separately by hour of arrest. We see that the number of overtime hours paid is relatively constant for non-felony arrests and increasing for felony arrests. Given that the average officer’s annual salary exclusive of overtime is approximately \$64,000, mean hourly pay is about \$33. Thus, given that overtime hours are paid out using 1.5 times an officer’s wage, a typical arrest that receives overtime is worth approximately \$153.

Having established that the probability of receiving overtime pay varies substantially over an officer’s shift, we next explore whether the frequency and quality of arrests vary. **Figure 2** plots the probability of an arrest (Panel a) and the probability of a misdemeanor or felony court conviction conditional on an arrest (Panel b) against the hour relative to the end of an officer’s shift. Referring to Panel (a), the probability of an arrest occurring in a given shift-hour is low, peaking at approximately 2.2 percent in the third hour of a shift. Interestingly, after rising over the first two hours of an officer’s shift, the probability of an arrest falls considerably — from approximately 2.2 percent at the -5 shift-hour to approximately 0.8 percent in the final hours of an officer’s shift, a decline of more than 60 percent. To address the possibility that the decline in arrests is due to officers processing prior arrests, we also plot the probability of an arrest in a given hour *conditional on having made an arrest previously* so that the magnitude of denominator falls throughout a shift. Given that the shape of the two curves in Panel (a) are strikingly similar, it is unlikely that the large decline in arrest activity that we observe throughout an officer’s shift is an artifact of incapacitation due to earlier shift enforcement activity.

If the number of arrests falls later in the shift, it is natural to consider whether the quality of arrests is also changing. A decline in the quality of arrests at the end of the shift would be consistent with the idea that officers reduce their evidentiary threshold for making late-shift arrests, perhaps in order to secure access to overtime pay.

These analyses can be thought of as a conceptual analog to the “hit rates” test for the presence of racial bias in some treatment (Knowles et al., 2001; Dominitz and Knowles, 2006; Persico and Todd, 2006). The presumption of the test is that, under no racial bias, the perceived guilt threshold for treatment should be the same across racial groups. In our setting, the presumption is that an officer who maintains the same standard for making an arrest throughout their shift should have the same success rate of an arrest in each shift-hour, where we measure success through court convictions and sentences.

Referring to panel (b) of Figure 2, we see that the probability that an arrest is sustained by a conviction *increases* throughout the work shift. Early shift arrests have a conviction rate of approximately 33 percent which increases to approximately 38 percent among arrests made at the end of an officer’s shift. In Section 5.1 we subject this analysis to further scrutiny using a series of controlled regressions which we describe in the following section. We also use a criminal sentence as an additional proxy for arrest quality.

day-of-week \times hour, division \times day-of-week \times shift and division \times month-year fixed effects.

4 Empirical Methods

4.1 Main Analysis

Our empirical strategy aims to identify the impact of an officer’s shift-hour (the number of remaining hours in the officer’s shift) on his or her propensity to make an arrest. In order to control for differences throughout the day in criminal behavior, we compare officers who work in the same community during the same hours but are at different points in their shift. This approach is possible because DPD officers work overlapping shifts within each sector to avoid significant disruptions during shift changes.¹¹ We provide a sense for the overlap in officer shifts in **Figure 3**. For each of the most common work shifts, the figure plots the average number of arrests for each hour in the day. While we plot only the most common shifts, we can see both that there are declines in arrest propensity within each shift and that there is substantial overlap in shifts throughout the day.

Our unit of observation is the officer-date-hour, which we denote by the triple ith . As discussed earlier, we include all days during an officer’s regular shift where they are assigned to a patrol unit. The primary outcome of interest Y_{ith} is an indicator variable for whether an arrest is made, and we run the following regression:

$$Y_{ith} = \sum_{k=-7}^{-1} \alpha_k \text{ShiftHour}_{ith}^k + X_{ith}\beta + u_{ith} \quad (1)$$

$$u_{ith} = \mu_i + \phi_{dwh} + \theta_{dws} + \xi_{dm(t)} + \epsilon_{ith} \quad (2)$$

Our controls X_{ith} comprise the log of volume of all calls and serious calls made within an officer’s sector in that hour. The coefficients of interest are α_k in (1) which tell us the relative probability of an arrest at each hour of a shift. For instance, $\alpha_{(-7)}$, the coefficient on the hour that is 7 hours from the end of an officer’s shift measures average arrest incidence relative to the first hour of an officer’s shift. Likewise, $\alpha_{(-1)}$ measures arrest incidence during the final hour of an officer’s shift.¹²

In order to identify shift-hour effects, we control extensively for potential confounders using a set of granular fixed effects which are presented in (2). We include interacted division \times day-of-week \times hour-of-day fixed effects, ϕ_{dwh} , in order to control for division-specific differences in arrest frequencies across all hours of the week. We account for additional aspects of the crime environment by controlling for interacted division \times year-month fixed effects, $\xi_{dm(t)}$, and division \times shift \times day-of-week fixed effects, θ_{dws} . These terms account for secular trends in the crime environment in each police division as well as persistent

¹¹Service disruptions can have large effects on clearance rates. See e.g., [Mastrobuoni \(2013\)](#).

¹²Since our models employ a large number of fixed effects, we use linear probability models in lieu of a non-linear estimator such as logit or probit. Estimates from a logit model with more aggregated fixed effects (year-month, division-hour, day of week) are substantively similar to our baseline estimates.

differences between shifts in each division with respect to arrest activity. These fixed effects are important insofar as dispatch may route officers in each shift to different sorts of service calls. Finally, all models condition on officer fixed effects, μ_i , to allow for time-invariant differences in arrest incidence across officers. We cluster standard errors at the officer and division-by-month level to account for arbitrary serial correlation in outcomes among arrests made by the same officer.

Two additional analyses merit description. First, we re-estimate (1) focusing on the probability that a 911 dispatch call leads to an arrest in a given shift-hour. Those models are estimated at the call-by-officer level; consequently, we additionally control for 911 call type fixed effects to account for differences in arrest probabilities for different types of emergency calls. We refer to our baseline sample as our “hourly sample,” and we refer to this sample of 911 calls as our “dispatch sample.”

Second, in order to test whether the probability of a criminal conviction and a sentence — both measures of arrest quality — differ throughout the work day, we let our unit of observation be an arrest-by-officer and regress an outcome of the arrest on the vector of shift-hour indicators, conditional on the same fixed effects described in (2).¹³

4.2 Effect of Off-Duty Work

By identifying the causal impact of an officer’s shift hour on his arrest behavior, equation (1) allows us to evaluate how the quantity and quality of arrests change throughout an officer’s shift. While these estimates provide a measure of officers’ relative preference or aversion to working overtime, they do not allow us to separately identify the importance of the monetary and non-monetary components of overtime work. Decomposing preferences into these two elements is the key to understanding how officer behavior responds to changes in financial incentives. We therefore seek to identify the impact of *changes* in the relative value of overtime work. Leveraging the fact that off-duty work changes the opportunity cost of an overtime spell, we test whether the shift-hour coefficients identified in (1) vary according to whether an officer is scheduled to perform off-duty work on a given day.

Our off-duty data include work that is both regularly occurring and idiosyncratic. Officers are required to notify their supervisor of their off-duty shift and receive approval beforehand. However, a 2018 audit of the department’s off-duty activity found that, in practice, officers regularly documented their shift and received approval after the occurrence of the shift (Smith, 2018). A resulting statistical concern is that, if processing an arrest requires that an officer work overtime and cancel a planned off-duty shift, that shift may never be officially reported. Such an occurrence may lead to a mechanical relationship between arrest activity and realized off-duty shifts even in the absence of a behavioral response.

¹³In our primary models, we do not condition on the arrest charge which is potentially endogenous, though we later estimate the model separately by arrest severity (violations, misdemeanors, and felonies) in order to explore mechanisms.

To avoid such a concern, we focus on an officer’s *regular* off-duty schedule. Specifically, we split an officer’s calendar into days of the week and quarters, and we say that an officer has a regularly planned off-duty shift on that day and quarter if more than 25 percent of those days have an off-duty spell.¹⁴ We then link these off-duty assignments to the officer’s police shifts to measure whether they have a regularly scheduled off-duty shift on that day. We say that an off-duty spell occurs before (after) the officer’s shift if some of the hours of the off-duty job occur before (after) the police shift and the off-duty shift ends before (after) the police shift ends (starts). All other spells, which comprise those that either fall entirely within a police shift or encompass it on both sides, are denoted as “during-shift” off-duty work.¹⁵

As in Equation 1, our unit of observation is an officer-date-hour. We regress our key outcome — a binary measure of whether an officer made at least one arrest — on an indicator variable for a predicted off-duty work spell. Notably we allow the effect of off-duty work to vary according to whether an arrest was made late (in the final four hours) versus early (in the first four hours) in an officer’s shift:

$$Y_{ith} = \sum_{q \in \{b,d,a\}} \gamma_E^q \widehat{OD}_{it}^q \times \text{Early}_{ith} + \sum_{q \in \{b,d,a\}} \gamma_L^q \widehat{OD}_{it}^q \times \text{Late}_{ith} \quad (3)$$

$$+ X_{ith}\beta + \mu_i + \phi_{dwh} + \theta_{dws} + \xi_{dm(t)} + \epsilon_{ith} \quad (4)$$

In (3), \widehat{OD}_{it}^q is equal to 1 if the officer worked an off-duty shift of type q on a given day of the week at least 25 percent of the time in a given quarter. Here, the superscript, q , refers to either pre-shift, during-shift, or post-shift off-duty work. We interact this term with an indicator for whether a shift-hour is late in an officer’s shift. We additionally control for officer, division-by-day of week-by-hour, division-by-shift and division-year-month fixed effects.

Notably, γ_E^q and γ_L^q are reduced form estimates of the relationship between arrests and the presence of scheduled off-duty work and, as such, are lower-bound estimates of the effect of realized off-duty work on arrest activity and overtime spells. In **Appendix Table 1**, we report first stage coefficients from a series of regressions of actual off-duty employment on predicted off-duty employment, predicting pre- and post-shift off-duty work using separate models. The entries along the diagonals report the relevant first stage coefficients while the off-diagonal elements establish that predicted pre-shift off-duty work is far less correlated with actual post-shift off-duty work and vice versa. For both before and after off-duty work,

¹⁴We take the start and end times for the shift to be the modal start and end time pair among the realized shifts on that day and quarter. If there is no modal start and end time pair, we designate the officer as not having a regular off-duty shift.

¹⁵The overlap between police shifts and off-duty shifts can occur because we are measuring regular rather than realized off-duty jobs, and an officer’s off-duty schedule may change in the middle of a quarter. In practice, during-shift off-duty work is rare in our data.

the estimated coefficients are approximately 0.6, indicating a strong correspondence between predicted and actual off-duty work.

4.3 Validity of Empirical Design

The key threat to identification in our estimation of equations (1) and (3) are unobserved differences in the criminal environment that may be correlated with our variables of interest. For our main regressions which identify the effect of shift-hour on arrest activity, one concern would be that, in periods where crime is high, supervisors may be more likely to have officers work on patrol rather than doing some other form of work. Therefore, certain hours may have a large share of officers working at the start of their shift and simultaneously experience high crime that generates high 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 partly address these concerns by flexibly controlling for differences in arrest propensity 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 and controls for the number of calls and serious calls in a division-hour address idiosyncratic changes in the environment over time. Similarly, our set of analyses where the unit of observation is a call taken by an officer also accounts for variation in criminal environment by allowing us to condition directly on the type of incident the officer is facing.

In order to probe the validity of our design, we conduct an imperfect but nevertheless informative balance test of our main variables of interest. Utilizing the same regressions described in equations (1) and (3), 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 shift-time predicts call volume. We present these results in **Appendix Table 2**. 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. If idiosyncratic variation in the composition of officer shift-times is uncorrelated with calls, the shift-hour coefficients should be jointly insignificant, which we test using an F -test. With respect to overall call volume, the 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 Appendix Table 2 test for a relationship between call volume and officer off-duty obligations. Similar to our first two columns, the coefficient on an off-

duty indicator tells us whether a division-hour with a higher than average share of officers working off-duty also has a higher than average call volume. 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.

Since the outcome variable is more aggregated than the variation we exploit in our design (which compares officers in the same division and time at different hours of their shift), these balance tests are an imperfect check on the validity of our design. However, the weak relationships reported in the table are consistent with the idea that the criminal environment faced by officer working in different shifts is close to being equivalent. Likewise, in **Section 5**, we present a series of robustness checks in which we consider various permutations of fixed effects, and our primary results are consistent across specifications.

5 Results

5.1 Main Results

We begin our discussion of the results by formally testing whether the frequency and quality of arrests vary over the course of an officer’s shift. Coefficients from equation (1) are plotted in **Figure 4**. In the figure, Panel (a) presents the arrest incidence regressions and Panels (b) and (c) present the conviction and sentence rate regressions, respectively. In each figure, the relevant shift-hour coefficient (on the dummy variable for each shift-hour) is plotted on the y -axis with the hour relative to the end of the shift plotted on the x -axis and 95 percent confidence intervals provide a boundary around the point estimates. In Panel (a), we see that the arrest probability initially rises by approximately 0.6 percentage points relative to the first hour arrest incidence of 1.3 percent, eventually rising to over 0.8 percentage points during the fourth hour of the shift before declining considerably. Consistent with the descriptive findings presented in Figure 2, arrests are 28 percent less frequent in the final hour of the shift than at mid-shift.

With respect to our measures of the quality of an arrest, we observe that the probability of a conviction and a sentence climb throughout the day, rising by approximately 9 percentage points each, relative to the first shift hour. For the probability of a conviction, a 9 percentage point change indicates a 27 percent increase in the conviction rate throughout the workday; for the probability of a sentence, this is a 39 percent increase. For both the probability of a conviction and a sentence, an F -test confirms that the probability in the final three hours of the shift is significantly higher than the probability in the first hour of an officer’s shift ($p < 0.05$).

Next, we consider the seriousness of arrest charges throughout an officer’s shift. If officers are strategically obtaining overtime wages by substituting low-level arrests at the end of their shift, we would expect to see the share of arrests for felonies decline during the final few hours

of an officer’s shift. Referring to panel (d) of Figure 4, we see little evidence for such an effect. While the shift-hour coefficients are statistically insignificant, we can rule out a decline in the felony share of arrests greater than five percentage points off a base of 16 percent share.

5.2 Robustness

In Section 5.1, we established that arrest activity declines towards the end of an officer’s shift and that officers tend to make higher quality arrests at the end of the workday. We argue that these effects are driven by officer aversion to working overtime. In this section, we consider alternative explanations and probe the robustness of our findings.

We begin by considering whether there are alternative explanations for the decline in arrest activity that we observe towards the end of an officer’s shift. A principal 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 concern, we ask whether 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. Models continue to condition on all of the standard fixed effects in equation (2) 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 22 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.¹⁶

A second concern is that there may be measurement error in the 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 mis-classified as having occurred either earlier in the workday or after the officer’s official workday has ended. Since

¹⁶In the model we present in Section 6, we estimate an officer’s probability of making an arrest in each hour of their shift and explicitly account for incapacitation from an earlier arrest. When we do so, we continue to find a reduction in arrest propensity throughout an officer’s shift.

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.¹⁷

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. As it turns out, the analysis in Figure 5 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 course of the plea bargaining process, could increase the probability of a conviction (Rehavi and Starr, 2014). In the left-hand panel of **Appendix Figure 3**, 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 charges after the first hour of the shift, the time path is roughly flat for the final seven hours of the shift, suggesting that the number of arrest charges does not substantially explain our court conviction finding.

Another explanation that could potentially rationalize our findings with respect to a court conviction 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 (Craig and Brown, 2014). In the right-hand panel of Appendix Figure 3, 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 5.4, the presence of an off-duty job after an officer’s

¹⁷We 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.

shift increases the opportunity cost of arrest processing. In Appendix Table 9, 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.

We also test the robustness of the results to alternative modeling strategies. In particular, given that the credibility of our estimates hinges on our ability to rule out competing explanations for the decline in arrest activity at the end of an officer’s shift using fixed effects, we re-estimate (1) using a variety of different sets of highly-granular fixed effects. In **Appendix Table 3**, we present tests for the decline in arrest propensity at the end of shift using various sets of fixed effects. For brevity, we replace the full set of shift-hour dummies with an F -statistic that tests the joint significance on the shift-hour coefficients for the final three hours of the officer’s shift. In the table, all models condition on police officer fixed effects but we vary the ways in which we account for the importance of place, work duties and time. In column (7) we specify a particularly saturated model in which we control for both interacted division \times date \times hour and division \times shift \times day-of-week fixed effects. Regardless of our choice of fixed effects, our principal estimates remain substantively similar and statistically significant.

5.3 Heterogeneity and Extensions

Having documented that our principal estimates are robust to a number of potential confounders and changes in functional form, we turn to exploring these potentially surprising results in greater detail. We begin by considering whether the incidence and quality of arrests varies according to the criminal seriousness of the arrest charge. **Figure 6** plots arrest incidence and the conviction rate by shift-hour separately for violations (low-level crimes that are considered by less criminally serious than misdemeanors) as well as misdemeanor and felony arrests. Estimates for all arrest types follow a similar pattern — there is an initial increase in the incidence of arrests early in an officer’s work shift followed by a decline during the final hours of the shift. However, the decline in arrest incidence is slightly larger for the least serious crimes — violations and misdemeanors — than for felonies.¹⁸ Such a result is consistent with the idea that police officers have far greater discretion over whether or not to make arrests for violations and misdemeanors than they do for felonies (Smith and Visser, 1981; Moskos, 2008a; Linn, 2009). Notably the results are inconsistent with the idea that police officers search for low-level arrests to make at the end of their shift. With respect to the conviction rate, we observe that the increase in the probability of a conviction over the course of a shift is driven almost entirely by violations and misdemeanor arrests suggesting that officers use their considerable discretion over whether to make arrests for less serious

¹⁸The felony results are nearly identical when we exclude felony drug charges.

crimes to moderate the degree to which they work overtime hours.¹⁹

Next, recognizing that recent literature finds evidence of racial disparities in policing (Goel et al., 2016; Fryer Jr, 2019; Goncalves and Mello, 2020), we consider whether officers make a larger number of late-shift arrests or lower quality late-shift arrests of minority citizens. We address this possibility in **Figure 7**. The left panels on top and bottom show the shift-hour effect on share Black and Hispanic arrestees, respectively. We find no evidence that the composition of arrestees becomes more Black or Hispanic near the end of the shift, though there is weak evidence that officers arrest relatively fewer Black individuals as their shifts progress. The right panels show the court conviction time paths separately by Black and Hispanic arrestees to test for whether the quality of arrests declines for either group. We find no difference in court convictions for Black arrestees throughout officers’ shifts, and the court conviction rate for Hispanic arrestees increases throughout the shift, consistent with the overall trend.

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 **Appendix Section A**, we investigate heterogeneity across officers in their late-shift arrest behavior, focusing on the officers with a disproportionately high share of arrests at the end of their shift. While a small minority of officers increase their arrest activity as the shift-end nears, their late-shift arrests are of similar composition and quality to their early-shift arrests.

5.4 Impacts of Off-Duty Work

Why don’t police officers exploit their broad discretion to make arrests in order to take advantage of overtime pay? We consider three possible drivers of these findings. First, productivity may decline throughout an officer’s shift due to declining ability throughout the workday (Vila et al., 2002; Vila, 2006).²⁰ Second, officers may place a high marginal value on their non-work time, such that they prefer to leave work rather than receive the additional pay of overtime (Mas and Pallais, 2019). Third, as an extension of the second case, officers may place a high marginal value on their non-work time, but their propensity to make late-shift arrests may be responsive to changes in overtime pay or other shocks to

¹⁹In **Appendix Table 4**, we probe the robustness of the characteristics of arrests throughout the shift, where we collapse shift-hour into a continuous variable for hours into the shift. Consistent with our baseline results, all models show an increase in court convictions and sentencing as the end of shift nears. With respect to the felony share of arrests, there is a small and sometimes-significant positive time path. Among violations and misdemeanor offenses, the share of arrests that yield a guilty conviction increases over time, while the time path of share guilty for felony arrests is noisy and insignificant.

²⁰While we could refer to this phenomenon as fatigue, a more precise conceptualization is a loss of ability to apprehend offenders that is independent of an officers preferences.

the marginal value of non-work time.

To distinguish between these possible cases, we exploit information on whether officers are working an off-duty shift on days in which they are also working a police shift. We use the presence of off-duty work prior to an officer’s police shift as an exogenous increase in the duration of the officer’s work day. If the observed decline in arrest propensity near the end of a shift is due to a decline in officer ability, then we should also observe a decline in arrest propensity on days in which an officer has pre-shift off-duty work. Conversely, the absence of an effect of pre-shift off-duty work on arrest propensity suggests that the observed end-of-shift decline in arrests is not the result of declining officer ability. To identify whether officers are responsive to the relative value of overtime pay, we exploit variation in post-shift off-duty work. When officers work an off-duty job after their scheduled work shift, a late-shift arrest compromises their ability to arrive on time for their off-duty job. We therefore interpret variation in post-shift off-duty work as leading to changes in the opportunity cost of making a late-shift arrest, which allows us to differentiate between the second and third stories above.

In **Table 3**, we document the frequency with which officers work in a secondary job. The average officer in our sample reports 45 off-duty shifts in a year, and the average length of a shift is 5.6 hours.²¹ On 9.8 percent of days with a police shift, officers also have an off-duty shift later in the day; this figure is 6.5 percent for off-duty shifts earlier in the day. As noted in Section 4.2, our empirical analysis will rely on indicators for an officer having a *regularly-scheduled* off-duty shift before or after their police work, which we define as days of the week where an officer reports an off-duty shift for at least 25 percent of dates in that quarter. On 6.9 percent of work days, officers have a regularly-scheduled off-duty shift *after* their work shift. Officers have a regularly scheduled off-duty shift *before* their work shift on 5 percent of work days. Consistent with Appendix Figure 1, we observe a reported pre-shift (post-shift) off-duty spell on 67 percent (69 percent) of days where we estimate an officer to have a regular off-duty spell; this number is 1.6 percent (3.2 percent) on non-regular days. These values indicate that our measure of regular off-duty work captures much of the variation in off-duty work.

Our estimates of the causal impact of off-duty work on arrest frequency, as per equation (3) are presented in **Table 4**. We present results for both the hourly sample in which the unit of observation is an officer-shift-hour, and the dispatch sample, where the unit of observation is each call taken by an officer. We likewise present estimates separately for felony and non-felony arrests. The first two rows provide an estimate of the impact of working an off-duty job prior to an officer’s police shift. We find only limited evidence that police officers make fewer late-shift arrests on days with pre-shift off-duty work. In the hourly sample, we see no evidence that arrests decline later in the officer’s shift; in the dispatch sample we estimate that late-shift arrests may have declined by 7 percent on days in which an officer is predicted

²¹A list of the most common off-duty jobs held by officers can be found in **Appendix Table 5**.

to work an off-duty shift prior to his work shift. Given that arrests decline by approximately 28 percent in the second half of the workday, the pre-shift off-duty effects indicate that ability loss explains as little as zero and, at most, 25 percent of the reduction in arrests that we observe at the end of the shift.²²

The second pair of rows in **Table 4** document the effect of off-duty work *after* an officer’s police shift. We see little evidence that early-shift arrests decline on days in which an officer is predicted to moonlight after his on-duty shift. However, late-shift arrests decline on these days by between 4 percent (in the overall sample) and 9 percent (in the dispatch sample).²³ Like the general reduction in late-shift arrest activity that we report in Figure 4, this effect is driven disproportionately by a decline in non-felony arrests. Notably, since these are reduced form estimates, they represent a lower bound on the impact of off-duty work — dividing these estimates by the first stage effect implies an estimate of between 7-15 percent. Overall, the evidence suggests that when the opportunity cost of overtime work rises, officers strategically reduce their arrest activity, focusing disproportionately on avoiding the most discretionary types of arrests. While these estimates indicate that officer behavior is sensitive to the presence of post-shift off-duty work, since fewer than 10 percent of shifts are followed by regular off-duty work, the scope for off-duty work to substantially distort the incentives to work overtime spells is modest. Taken as a whole, off-duty work likely explains only a small share of the 28 percent reduction in arrest activity towards the end of an officer’s workday.

The final column of the table shows an estimate of the impact of off-duty work on whether an officer receives overtime on a given day. The unit of observation for this regression is an officer-shift, and we include all of our fixed effects from equation (4) with the exception of division \times day-of-week \times hour. We find that both pre-shift and post-shift off-duty work lead to significant reductions in the likelihood of working overtime. Relative to a baseline mean of 28.3 percent, pre-shift and post-shift off-duty work lead to 3.4 and 5.4 percentage point reductions in the likelihood of overtime, respectively. These impacts are larger than the arrest impacts. Officers can receive overtime for several forms of activity, not solely arrests, and these estimates suggest that officers also reduce their use of general overtime that is not a function of making an arrest on days in which they are employed in an off-duty job.

While the results above indicate that officers reduce their late-shift arrest activity on days with post-shift off-duty work, a natural question is what happens to the quality of arrests that are made. Since our baseline findings suggest that arrest declines are accompanied by

²²We explore the robustness of our findings in **Appendix Table 6** which presents estimates arising from the same mix of fixed effects approaches used to test the robustness of our principal findings. The effect of post-shift off-duty work is robust to specification and there is little evidence for an effect of pre-shift off-duty work in any of the models. In **Appendix Table 7** we present separate estimates for short (≤ 4 hours) versus long (> 4 hours) pre-shift off-duty work spells and find that the late-shift arrest reduction is specific to shifts with a long post-shift job.

²³The effect is especially strong for regular off-duty work that occurs *directly* after an officer’s shift rather than later in the day — we present these results in **Appendix Table 8**.

increases in arrest quality, we should expect that arrest quality is higher on days with off-duty work. The results of such an analysis are presented in **Appendix Table 9**. We use the same regression of Equation (3), where the unit of observation is an arrest, and each column considers a different outcome. We do not find evidence that arrests made when officers are working off-duty are more likely to lead to a guilty conviction or a prison sentence, nor are they significantly more likely to be a felony or a misdemeanor/felony. While these results may be surprising given our previous estimates, our standard errors are likely too wide to make any definitive conclusions. In Section 6, we show that the expected increase in guilty convictions for arrests on post-shift off-duty days, though positive, is small and within the confidence intervals we present in Appendix Table 9.

Our estimates of the off-duty impacts suggest that the decline in officer arrest activity throughout their shift cannot be explained by declining ability, since we find small and inconsistently significant estimates for the impact of pre-shift off-duty work. Instead, we argue that the decline is more consistent with an aversion to working overtime. Our finding that post-shift off-duty work also leads to a decline suggests that officers are responsive to the opportunity cost of arrests. The following section synthesizes these findings with our baseline results in a simple dynamic model. By doing so, we can identify the relative weight officers place on their arrest activity and their personal overtime considerations.

6 Model of Officer Arrest Decisions

We have shown that officers reduce their arrest activity near the end of their shifts and that the quality of the arrests made — as proxied by court convictions and sentencing — increases as the end of a shift nears. We then showed that officers who work an off-duty shift after their police shift exhibit a further decline in late-shift arrests. In this section, we present a simple dynamic model of an officer’s arrest decisions. We match this model to our empirical results, allowing us to estimate officer preferences over arrests of guilty individuals and avoidance of arrests of innocent individuals (i.e., their pro-social motivations), their relative value for working overtime, and the value they place on off-duty and overtime pay. This model will allow us to discern officer corruptibility and consider the impact of changes to financial incentives on the quantity and quality of arrests.

6.1 Setup

Guilt Signal Structure: An officer works a shift with $T = 8$ periods. In each period $t = 1, \dots, 8$, the officer encounters a potential arrestee, whose guilt is governed by probability, ρ . We designate whether an individual is guilty with the indicator $g \in \{0, 1\}$. The officer does not observe the individual’s guilt but instead observes a noisy signal S that is correlated

with guilt in the following manner:

$$(S|g) \sim N(g\mu, 1)$$

The mean of the signal distribution for guilty individuals, μ , dictates how well an officer is able to differentiate signals between guilty and innocent individuals. Using this signal and his or her prior on the probability of guilt, the officer generates an updated probability of guilt for the individual, $\tilde{p}(s)$:

$$\tilde{p}(s) = \frac{\phi(s - \mu)\rho}{\phi(s)(1 - \rho) + \phi(s - \mu)\rho}$$

Officer Per-Period Objective Function: In a given period, an officer has the choice of arresting the individual he or she encounters, and the contemporaneous value that the officer derives from his or her decision is a weighted sum of the value of correctly arresting a guilty individual and correctly not arresting an innocent individual:

$$v(a, s) = a(1 - \lambda)\tilde{p}(s) + (1 - a)\lambda(1 - \tilde{p}(s))$$

λ is a parameter which dictates the importance that an officer places on avoiding arrests of innocent individuals relative to arresting guilty individuals. Henceforth we refer to this parameter as the officer's risk preference.

If an officer were choosing arrest activity in a static setting, where the above equation is his sole objective function, he or she would arrest an individual if their posterior guilt probability is $\tilde{p}(s) \geq \lambda$. This formulation of the officer's decision problem has been used previously by Prendergast (2003, 2007) to model public sector agents and Alesina and La Ferrara (2014), Arnold et al. (2018), and Arnold et al. (2020) for modeling judge behavior. Our model uses this objective functions as its base and adds a dynamic component.

If the officer makes an arrest, there is a probability ϕ that the officer may need to process the arrest, during which time the officer is not patrolling and is unable to arrest any additional suspects. If the officer does process the arrest, then every period thereafter there is a probability p that the officer must continue to process, until the officer is no longer processing. The model structure is depicted visually in **Appendix Figure 6**.

When processing the arrest, the officer cannot observe individuals on patrol. Therefore, the expected value they receive is the baseline expected value from an innocent individual not being arrested:

$$v^{np} = \lambda(1 - \rho)$$

End of Shift: After the final period of the day, if the officer is still processing an arrest, he receives a value of being at work and receiving overtime, $V_{ot} = c_{ot} + b \cdot \text{Pay}_{OT}$, where Pay_{OT} denotes the expected overtime pay received by the officer. If the officer is not processing an arrest at the end of the shift, he receives value, $V_0 = [c_{od} + b \cdot \text{Pay}_{OD}] \cdot \text{OD}$, where OD is an indicator for whether the officer is working off-duty after their police shift, and Pay_{OD} is

their pay for that off-duty shift. The coefficients in front of overtime and off-duty pay are the same, so that we assume the officer values a dollar of pay equally from both sources.

Note that we have normalized the value of ending the day without working overtime or working off-duty to 0, so the parameters c_{ot} and c_{od} indicate the costs or benefits to working relative to going home that cannot be explained by pay. For example, one potential cost to missing an off-duty shift is that the officer may be fired from their secondary job, leading to a greater income loss than the pay of the single missed shift. Our specification assumes that this kind of loss is adequately captured by c_{od} , so that the potential future earnings loss is the same regardless of the pay of a single shift.²⁴

Value Functions: All the utility values presented so far have been single-period, and here we present value functions that represent an officer's utility for the current period and the expectation of utility in later periods. The value function of an officer who is patrolling in period t , conditional on their arrest choice and guilt signal, is denoted by $V_t^p(a, s)$. When indicating value of patrolling prior to conditioning on arrest choice and signal, we use the variable $V_t^p \equiv E_s[\max_a V_t^p(a, s)]$.

The value function for an officer in the final period T is a direct function of the values of working overtime and leaving work:

$$V_T^p(a, s) = v^p(a, s) + a\{\phi \cdot V_{ot} + (1 - \phi) \cdot V_0\} + (1 - a)V_0$$

An officer's value function from making and not making an arrest decision in period $t < T$ are, respectively,

$$\begin{aligned} V_t^p(a = 1, s) &= v^p(a = 1, s) + \phi \cdot V_{t+1}^{np} + (1 - \phi) \cdot V_{t+1}^p \\ V_t^p(a = 0, s) &= v^p(a = 0, s) + V_{t+1}^p \end{aligned}$$

where V_t^{np} indicates the value of not patrolling in period t ,

$$V_t^{np} = v^{np} + p \cdot V_{t+1}^{np} + (1 - p) \cdot V_{t+1}^p$$

We add two additional features to the model to better fit the data. First, we impose that officers begin on patrol with probability p_{start} . If an officer begins off of patrol, they have probability $1 - p$ to enter patrol each period. We add this probability to match that officers have a lower arrest rate earlier in the shift than in the middle. Second, we allow an exogenous probability p_{ot} that an officer receives overtime regardless of processing an arrest at the end of the shift. We do so to match that officers sometimes receive overtime on days where they make no arrests.

Note that arrests only enter the officer's utility function through the value of arresting a guilty individual and the value of avoiding arresting an innocent individual. Our interpretation of these expressions are as indicators of intrinsic value officers place on these outcomes

²⁴From speaking with officers in the department, it appears to be relatively common to find a fellow officer to take over a missed shift, so the likelihood of being fired from a singled missed shift seems small.

(Bénabou and Tirole, 2006; Besley and Ghatak, 2018). One concern with this interpretation is the possibility that officers may receive some workplace benefit from achieving these outcomes. If that is the case, officers may appear to intrinsically value arrests but actually value another outcome that they achieve through arrests, such as career advancement or higher pay. One strength of our setting is that officer incentives are quite low-powered. Promotions from officer to corporal, then sergeant and lieutenant, are allocated through promotion exams. Individuals are ranked by their exam performance and promoted as positions open, with no reference to an officer’s number and quality of arrests.²⁵ Therefore, career advancement for most officers is not a function of arrest activity. However, some workplace benefits can be derived from arrests. For example, supervisors have discretion in providing non-arrest overtime to officers and may favor officers who make many arrests that lead to conviction. Further, officers may feel pressure from peers to make a sufficient number of high-quality arrests (Bandiera et al., 2009, 2010). Therefore, we interpret the officer’s utility for arrests as reflecting intrinsic utility, with the important caveat that our estimates may also partly capture low-powered workplace incentives and social pressure over arrest activity.

6.2 Solution

The optimal solution for the officer will be characterized by a series of threshold values, s_t^* , for each period, which indicate that individuals whose signals are above s_t^* will be arrested, and those below will not be arrested. The threshold values have the following solution:

$$\tilde{p}(s_t^*) = \lambda + \phi[V_{t+1}^p - V_{t+1}^{np}] \quad (5)$$

This equation offers a simple interpretation for the officer’s rule. When the value of staying on patrol is higher than the value of being off of patrol in the following period, the officer raises his guilt threshold. Note that if an officer were simply maximizing the within-hour value, $v(a, s)$, his threshold rule would be $\tilde{p}(s_t^*) = \lambda$ for every period. And if it were the case in the dynamic setting that $\phi = 0$, the threshold rule would also be the same in all periods and equal to λ . Therefore, the distortion in officer behavior relative to the static case comes entirely from the fact that the arrest decision affects whether he or she is on patrol in the following period. In the final period of the day, where $V_{t+1}^p = V_0$ and $V_{t+1}^{np} = V_{ot}$, the distortion is based on the officer’s relative taste for working late, which may depend on whether the officer has off-duty work scheduled. However, in earlier periods there may be some additional deviation from λ because of the opportunity cost of not being able to arrest guilty individuals in future periods.

²⁵For higher ranks of the department, promotion is discretionary and can be dependent on any measure of job performance.

6.3 Model Estimation

To estimate the model, we match four sets of moments from the data. First, the model uses the probability of making an arrest in each hour t . These are matched by the probability that an officer is working on patrol in hour t times the probability that an individual’s guilt signal, s , is greater than s_t^* . Second, we match the probability of overtime conditional on an arrest as well as the probability of overtime when making no arrest. We will match these by calculating that an officer is processing an arrest at the end of the shift conditional on making an arrest in each hour.²⁶ Third, we match the probability of conviction conditional on arrest. For the purposes of estimating the model, we treat a court conviction as “ground-truth” guilt; the absence of a conviction indicates that the individual is innocent. Another interpretation of this assumption is that we are assuming that the officer’s notion of guilt is based on whether the individual can be convicted in court. Therefore, the probability of conviction conditional on arrest is equal to the probability of an individual’s guilt conditional on having $s > s_t^*$.

To estimate the parameters of the officers’ value of overtime and off-duty, we use the reduced form arrest impacts of having an off-duty shift after their police shift. Note that the pay levels Pay_{OT} and Pay_{OD} but also the level preferences c_{ot} and c_{od} differ between overtime and off-duty work. We therefore need some variation in payment within off-duty spells to separately identify both parameters. To do so, we run a version of Table 4 where we split off-duty shifts into four hours or shorter and longer than four hours, while maintaining the choice from our main specification of splitting the effects into the first and last four hours of an officer’s shift. The results of this regression are presented in **Appendix Table A7**, which provides our fourth set of moments. We match the effects of short and long after-shift off-duty spells on the first and last four hours of an officer’s shift, giving us four moments.

The model also requires that we choose values for Pay_{OT} and Pay_{OD} . For overtime pay, we take the average officer’s salary in our sample, \$64,443, and the average number of overtime hours worked, 2.72, multiplied by 1.5, and get an average payment of \$131.8. For off-duty work, we observe that short spells (\leq four hours) have an average duration of 3.4 hours, and long spells ($>$ four hours) have an average duration of 6.4 hours. While we do not observe the hourly pay of officers for their off-duty jobs, the available figures online indicate an hourly rate of \$30-40.²⁷ We choose a uniform value of \$35 and impute average payments of \$119.3 and \$224.3 for short and long off-duty shifts, respectively.

For the first three sets of empirical moments, we use estimates based on the sample of shifts where an officer is not working off-duty, and the moments are matched to the model with $OD = 0$. We then estimate Equation 1 with our baseline set of fixed effects. Our moment for each hour and outcome is $\hat{\mu}_8 + \hat{\theta}_k$, where $\hat{\mu}_8$ is the average for the first hour of

²⁶Note that an officer can make an arrest in a certain hour and be processing a *later* arrest when he receives overtime, which is accounted for in constructing these moments from the model.

²⁷<https://smallbusiness.costhelper.com/security-guard.html>

the shift and $\hat{\theta}_k$ is the estimated coefficient for each hour.

In total, we have ten parameters to be estimated and 29 moments (three moments times eight shift-hours plus overtime probability when not arresting and four off-duty moments) with which to estimate the model. Despite the over-identification of the model, we show that the fit of all the moments is quite good.

6.4 Model Estimates and Counterfactuals

Model estimates are presented in **Table 5**. We estimate that in any given hour, the probability that an officer encounters a guilty individual is 0.037. The officer’s average ability μ is 1.555, indicating that the distribution of signals for innocent and guilty individuals is quite overlapping. This low value explains why the true guilt rate is estimated to be significantly higher than any arrest rates we observe throughout the shift. We find that officer’s trade-off between releasing innocent individuals and arresting guilty individuals, λ , is 0.154, meaning that, in a static setting, officers would arrest individuals with a posterior probability of guilt higher than 15 percent. When an arrest is made, the probability of the officer processing the arrest is 0.791, and conditional on processing, the probability of continued processing in each successive period is 0.761.

The value of \$1 of either overtime or off-duty income is $\hat{b} = 0.000111$. We can use this estimate to provide a dollar value for an officer’s pro-social motivations. Dividing $\hat{\lambda}$ by \hat{b} yields an estimate of the value, in dollars, of correctly failing to arrest an innocent individual. This estimate is \$1,382. Likewise, deflating $1 - \hat{\lambda}$ by this number yields an estimate of the value that officers placed on correctly arresting a guilty individual — \$7,621. These numbers are large, especially in comparison to the daily salary rate of \$257. As such, the model suggests that officers place great value on the avoidance of Type I and especially Type II errors.

One valuable way to interpret our estimates for officer pro-social motivation is to consider the degree of financial incentives needed to distort their activity in un-ambiguously harmful ways. We consider here two measures of *corruptibility*. First, our estimates imply that an officer in the last hour of her shift would need to expect at least \$3,548 in overtime pay to make an arrest of an individual she knows with certainty is innocent.²⁸ Conversely, an officer in her last hour of the shift would need to have at least \$8,024 of income waiting at an off-duty job to forgo arresting an individual she knows with certainty is guilty.^{29,30}

Since we do not observe overtime or off-duty pay of this magnitude, these figures neces-

²⁸We solve for the value of Pay_{OT} such that $\phi[c_{OT} + b \cdot \text{Pay}_{OT}] = \lambda$.

²⁹We solve for the value of Pay_{OD} such that $\phi[c_{OD} + b \cdot \text{Pay}_{OD}] = (1 - \lambda) + \phi[c_{OT} + b \cdot \text{Pay}_{OT}]$, where we use our baseline value of Pay_{OT} .

³⁰Note that our estimates indicate that officers care more about Type II errors than Type I errors. These estimates follow from the fact that $\hat{\lambda}$ is lower than 0.5. However, this preference is tempered by the fact that officers’ guilt threshold is adjusted *upwards* due to an aversion to working overtime, as can be seen in Figure 9A.

sarily require extrapolation from the smaller degree of variation in off-duty pay we used to estimate the model. Consequently, we think of these estimates as lower bounds for the true corruptibility values. For example, if an officer is willing to reduce their guilt threshold by 5% when they receive a higher value of overtime pay, they will likely need to receive more than ten times that increase to reduce their threshold by 50%, though our estimates assume a linear extrapolation.

In **Figure 9A** and **Figure 9B**, we document the fit of the model to our moments. The estimates fit the empirical moments well, especially when noting the relative parsimony of the parameters relative to the number of moments. We are able to capture the increase and then decline in arrest activity throughout the shift, the increase in court convictions throughout the shift, and the increase in overtime throughout the shift. The model estimate for the time path of overtime rates is positive and convex, driven by our assumption that the processing rate is governed by an hourly rate of continued processing. In contrast, the empirical time path is quite linear, which leads to a slight deviation of model fit from the matched moments.

The bottom right panel of Figure 8A shows our estimate for the hourly probabilities of guilt above which an officer would arrest a suspect. Consistent with the findings from our reduced form estimates, we find an increase in the arrest threshold as the officer advances in his shift, confirming that the decline in arrest propensity is driven by a change in arrest behavior rather than just a gradual incapacitation from prior arrests. Note also that the guilt threshold throughout the entire shift is above the hypothetical threshold from a static version of the officer's problem, indicating that officer arrest propensity is impacted by an aversion to working overtime at all hours of the shift.

Figure 8B documents the fit of the model to the off-duty effects. All of our model-simulated moments are within the empirical confidence intervals, suggesting a good fit. We also capture the fact that the point estimate are more negative for late-shift arrest effects than early-shift. However, the model appears to simulate slightly more negative effects for early-shift arrests than our empirical moments. This discrepancy is due to the fact that declines in arrests must be driven solely by an increased aversion to overtime work, and the high rate of processing and continued processing leads the model to conclude that even *early-shift* arrests must be reduced. While this finding highlights a slight deviation of the model from the empirical moments, it also documents how the model's predictions and its fit to the data are non-trivial.

As we note in Section 5.5, we do not find an increase in convictions for arrests made on off-duty days, despite the intuition that declining arrests should lead to higher-quality arrests. While we do not include these estimated effects as moments in our model estimation, we simulate what our model would predict for the off-duty conviction effects, and we compare them to our empirical findings in **Appendix Figure 7**. While the simulated model effects are positive (and the empirical effects are negative and insignificant), the estimates are

comfortably within the confidence intervals of the empirical estimates. This finding indicates to us that, though the model and empirical estimates are roughly similar, our sample size limits our ability to precisely estimate these impacts.

A crucial question we can address with this model is how an officer’s arrest behavior would change with an adjustment to the value of working overtime, which we present in **Figure 9**. The red line presents the model fit to our empirical moments. We first consider a tripling of the pay for overtime work, corresponding to the blue line. As expected, we observe an increase in arrest propensity and decline in guilt convictions throughout the shift.³¹ Interestingly, these impacts are relatively small. The elasticity of arrest probability to overtime pay is 0.03, indicating that a 100% increase in overtime pay would lead to only a 3% increase in arrest propensity. This finding is due to both our assumption that officers value overtime pay the same as off-duty pay and to our finding that off-duty work does not lead to large declines in arrests.

A similar counterfactual we consider is how an officer’s arrest behavior would appear if he did not exhibit a preference for not working overtime. We simulate this possibility by changing the value of V_{ot} to 0, which we plot in the green line. Surprisingly, this officer exhibits an *increase* in arrest activity throughout the shift and a corresponding decline in court convictions. Intuition might lead us to believe that the overtime-indifferent officer would have a constant arrest rate and guilt threshold throughout the shift. However, arrests in the earlier hours impose the possibility of processing during the shift and precluding future arrests. What is especially striking about this finding is that, without estimation of the dynamic model above, an observed arrest profile that increases throughout the shift would likely be taken as evidence in favor of the “collars for dollars” story, namely that officers reduce their arrest quality in order to receive overtime pay. Our model indicates that, on the contrary, a completely altruistic officer may exhibit an increasing arrest rate and decreasing guilt threshold across their shift because of the higher opportunity cost earlier in the shift.

6.5 Alternative Modeling Choices

The model we present is quite simple, and we naturally omit some dimensions of the officer’s decision framework. We consider here two additional complications and present additional results in **Appendix Table 10**. The first column presents our baseline model parameter estimates and subsequent columns present alternative specifications.

First, our baseline model imposes that an individual is convicted in court if and only if he is truly guilty. However, in reality, at least some innocent individuals are convicted and

³¹The overall guilt conviction impact is in principle ambiguous: while convictions in all hours are now less likely, the arrest composition more heavily weights late-shift arrests. The overall court conviction rate is lower with triple overtime pay, indicating that the former effect outweighs the latter.

at least some guilty individuals are not convicted. Letting true guilt be denoted by g and a conviction be denoted by c , we estimate a version of the model in which we suppose that errors occur in 10 percent of cases and are symmetric: $Pr(c = 1|g = 0) = Pr(c = 0|g = 1) = 0.1$. With this additional feature, the court conviction rate for an hour, c_h , and true guilt rate g_h , are related by $c_h = 0.9 * g_h + 0.1 * (1 - g_h)$.

The result of this alternative specification is presented in the second column of Appendix Table 10. The most notable change relative to the baseline estimates is that λ , indicating the relative importance of accurate arrests and non-arrests, declines from 0.154 to 0.086, suggesting that officers place an even higher importance on correct arrests relative to non-arrests. The value of b also increases slightly. These changes ultimately lead to a decline in values of an arrest from \$7,621 in the baseline model to \$7,516, and a non-arrest from \$1,382 in the baseline model to \$703. While these changes are notable, the general magnitude of values is similar, and the expected value of an hour of patrol relative to non-patrol is not substantially changed, \$58 to \$62.

Another change we consider is to replace our estimates of the effect of off-duty work on arrest activity, which can be thought of as “reduced-form” estimates, with estimates scaled by $1/0.57$, the first-stage coefficient from Appendix Table A.1. The resulting estimates are presented in column 3 of Appendix Table 10. The value of λ is essentially unchanged relative to the baseline model. However, b has increased substantially, from 0.00011 to 0.00020. This increase leads to a similar decline in intrinsic value for arrests as our court error model. The value of a non-arrest of an innocent individual declines from \$1,382 to \$761.4, and the value of an arrest of a guilty individual declines from \$7,621 to \$4,193. The final column presents both adjustments together, and the values change to \$384.7 and \$4,111.2, respectively.

7 Conclusion

In this paper, we evaluate how officers trade off their intrinsic motivation over workplace outcomes with their private preferences for overtime work. In particular, we ask whether the quantity and quality of arrests are substantially distorted by officer financial incentives. Using unique administrative data which link records on 911 calls, police officer shift assignments, off-duty work, arrests, and associated court outcomes in Dallas, Texas, we find that officers significantly reduce their frequency of arrests as their shift-end nears. This finding is not explained by officers being taken out of circulation having made earlier shift arrests, nor is it explained by “arrest trading” between officers or a formal or informal policy to route officers to fewer calls at the end of their shift. We next find that the conviction and sentencing rates for arrests increase at the end of the shift. Like Chan (2018), who studied shift work among emergency department physicians, our findings are consistent with an aversion to working overtime. Leveraging variation in off-duty work, which causes a shift in the opportunity cost of overtime work, we also find that officers reduce their last-shift arrests

when the cost of making an arrest is highest. Feeding our estimates into a dynamic model of officer arrest decisions, we estimate very high intrinsic values by the officers on arresting guilty individuals and not arresting innocent individuals, suggesting that the average officer exhibits a high degree of pro-social motivation in their job.

Our study suggests several avenues for future research. While our findings run contrary to prevailing beliefs about arrests and overtime pay, we focus on a single police department. As such, the results we report are inevitably shaped by the institutional practices of the Dallas police department. In addition to creating a common set of expectations regarding productivity, our sample focuses on officers working eight-hour shifts. As research suggests that officers' preferences over overtime work can be sensitive to shift length (Amendola et al., 2011), it is possible that the late-shift decline in officer productivity that we observe could be even larger in a department that uses 10 or even 12 hour shifts. Of course, addressing this question will require similarly rich data available across several departments each of which, like Dallas, use overlapping shifts.

Second, our estimates of the pro-social motivation of officers are estimated from a relatively fixed group of individuals who face a constant pay schedule and employment contract. How would the pro-social motivation of the police force respond to a change to the employment contract? A large share of the literature on public sector workers focuses on the labor supply of agents and how pro-social motivation is affected by changes to job characteristics (Fisman et al., 2015; Ashraf et al., 2020). Connecting our approach to an investigation of police labor supply would be valuable for evaluating any policy changes related to workplace characteristics.

Finally, how surprising are our estimates of officers' intrinsic valuations for arrests and non-arrests? What value would the average civilian place on these outcomes, and what is the socially optimal degree of pro-social motivation? Our study is among the first to provide evidence of pro-social motivations identified directly from changes to the opportunity cost of a workplace activity, so there is for now an absence of other estimates to consult as comparison. While there is a large literature on the social value of crime prevention (Aldy and Viscusi, 2008; Cohen and Piquero, 2009; Chalfin, 2015), there is hardly any research on the value individuals place on the arrest of a criminal suspect after a crime has been committed. Future research on estimates of pro-social motivation of public sector workers and on estimates of civilians' valuations of arrests will provide valuable guidance in interpreting our finding that officers appear to place high valuations on their workplace outcomes.

References

- Aldy, J. E. and W. K. Viscusi (2008). Adjusting the value of a statistical life for age and cohort effects. *The Review of Economics and Statistics* 90(3), 573–581.
- Alesina, A. and E. La Ferrara (2014). A test of racial bias in capital sentencing. *American Economic Review* 104(11), 3397–3433.
- Amendola, K. L., D. Weisburd, E. Hamilton, G. Jones, M. Slipka, A. Heitmann, et al. (2011). *The shift length experiment: What we know about 8-, 10-, and 12-hour shifts in policing*. Police Foundation Washington, DC.
- Anechiarico, F. and J. B. Jacobs (1994). Visions of corruption control and the evolution of american public administration.
- Arnold, D., W. Dobbie, and C. S. Yang (2018). Racial bias in bail decisions. *The Quarterly Journal of Economics* 133(4), 1885–1932.
- Arnold, D., W. S. Dobbie, and P. Hull (2020). Measuring racial discrimination in bail decisions. Technical report, National Bureau of Economic Research.
- Ash, E. and W. B. MacLeod (2015). Intrinsic motivation in public service: Theory and evidence from state supreme courts. *The Journal of Law and Economics* 58(4), 863–913.
- Ashraf, N., O. Bandiera, E. Davenport, and S. S. Lee (2020). Losing prosociality in the quest for talent? sorting, selection, and productivity in the delivery of public services. *The American Economic Review* 110(5), 1355–94.
- Ater, I., Y. Givati, and O. Rigbi (2014). Organizational structure, police activity and crime. *Journal of Public Economics* 115, 62–71.
- Baicker, K. and M. Jacobson (2007). Finders keepers: Forfeiture laws, policing incentives, and local budgets. *Journal of Public Economics* 91(11-12), 2113–2136.
- Bandiera, O., I. Barankay, and I. Rasul (2009). Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica* 77(4), 1047–1094.
- Bandiera, O., I. Barankay, and I. Rasul (2010). Social incentives in the workplace. *The Review of Economic Studies* 77(2), 417–458.
- Banuri, S. and P. Keefer (2016). Pro-social motivation, effort and the call to public service. *European Economic Review* 83, 139–164.
- Bénabou, R. and J. Tirole (2006). Incentives and prosocial behavior. *The American Economic Review* 96(5), 1652–1678.
- Benson, B. L., D. W. Rasmussen, and D. L. Sollars (1995). Police bureaucracies, their incentives, and the war on drugs. *Public Choice* 83(1-2), 21–45.

- Bertrand, M., R. Burgess, A. Chawla, and G. Xu (2020). The glittering prizes: Career incentives and bureaucrat performance. *The Review of Economic Studies* 87(2), 626–655.
- Besley, T. and M. Ghatak (2018). Prosocial motivation and incentives. *Annual Review of Economics* 10, 411–438.
- Blumenson, E. and E. Nilsen (1998). Policing for profit: The drug war’s hidden economic agenda. *The University of Chicago Law Review* 65(1), 35–114.
- Bowles, R. and N. Garoupa (1997). Casual police corruption and the economics of crime. *International Review of Law and Economics* 17(1), 75–87.
- Brewer, R. C. and H. D. Gilliam (1979). Production appraisal of police patrol. *American Journal of Criminal Justice* 4(2), 55–62.
- Bronstein, N. (2014). Police management and quotas: Governance in the compstat era. *Columbia Journal of Law & Social Problems* 48, 543.
- Cassell, P. G. (2018). Overstating america’s wrongful conviction rate: Reassessing the conventional wisdom about the prevalence of wrongful convictions. *Arizona Law Review* 60, 815.
- Cavallaro, R. (1998). Police and thieves.
- Chalfin, A. (2015). Economic costs of crime. *The Encyclopedia of Crime and Punishment*, 1–12.
- Chan, D. C. (2018). The efficiency of slacking off: Evidence from the emergency department. *Econometrica* 86(3), 997–1030.
- Cohen, M. A. and A. R. Piquero (2009). New evidence on the monetary value of saving a high risk youth. *Journal of Quantitative Criminology* 25(1), 25–49.
- Craig, L. and J. E. Brown (2014). Weekend work and leisure time with family and friends: Who misses out? *Journal of Marriage and Family* 76(4), 710–727.
- Department of Justice, U. (2015). *The Ferguson Report: Department of Justice Investigation of the Ferguson Police Department*.
- Dixit, A. (1997). Power of incentives in private versus public organizations. *The American Economic Review* 87(2), 378–382.
- Dixit, A. (2002). Incentives and organizations in the public sector: An interpretative review. *Journal of Human Resources*, 696–727.
- Dominitz, J. and J. Knowles (2006). Crime minimisation and racial bias: what can we learn from police search data? *The Economic Journal* 116(515), F368–F384.
- Farber, H. S. (2015). Why you cant find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics* 130(4), 1975–2026.

- Feuer, A. and J. Goldstein (2018). The arrest was a bust. the officers got overtime anyway.
- Finan, F., B. A. Olken, and R. Pande (2017). The personnel economics of the developing state. In *Handbook of Economic Field Experiments*, Volume 2, pp. 467–514. Elsevier.
- Fisman, R., N. A. Harmon, E. Kamenica, and I. Munk (2015). Labor supply of politicians. *Journal of the European Economic Association* 13(5), 871–905.
- Forst, B. (1982). *Arrest convictability as a measure of police performance*, Volume 73. US Department of Justice, National Institute of Justice.
- Fryer Jr, R. G. (2019). An empirical analysis of racial differences in police use of force. *Journal of Political Economy* 127(3), 1210–1261.
- Fyfe, J. J. (1984). Testing” the limits of law enforcement”.
- Getter, L., J. Leen, and G. Epstein (2014). Innocent caught in web of cops’ overtime abuse.
- Glaeser, E. L. and C. Goldin (2006). Corruption and reform: Introduction. In *Corruption and Reform: Lessons from America’s Economic History*, pp. 2–22. University of Chicago Press.
- Goel, S., J. M. Rao, R. Shroff, et al. (2016). Precinct or prejudice? understanding racial disparities in new york citys stop-and-frisk policy. *The Annals of Applied Statistics* 10(1), 365–394.
- Goldstein, H. (1963). Police discretion: The ideal versus the real. *Public Administration Review*, 140–148.
- Goncalves, F. and S. Mello (2020). A few bad apples? racial bias in policing. *Available at SSRN 3627809*.
- Greene, J. R. (2000). Community policing in america: Changing the nature, structure, and function of the police. *Criminal Justice* 3(3), 299–370.
- Gregg, P., P. A. Grout, A. Ratcliffe, S. Smith, and F. Windmeijer (2011). How important is pro-social behaviour in the delivery of public services? *Journal of Public Economics* 95(7-8), 758–766.
- Groeneveld, R. F. (2005). *Arrest discretion of police officers: The impact of varying organizational structures*. LFB Scholarly Pub.
- Heckman, J. J., J. A. Smith, and C. Taber (1996). What do bureaucrats do? the effects of performance standards and bureaucratic preferences on acceptance into the jtpa program. Technical report, National Bureau of Economic Research.
- Holmstrom, B. and P. Milgrom (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics & Organization* 7, 24.

- Kaplan, J. (2019). Uniform crime reporting (ucr) program data: County-level detailed arrest and offense data.
- Kelling, G. L. (1999). *Broken windows and police discretion*. US Department of Justice, Office of Justice Programs, National Institute of Justice.
- Kelly, B. D. and M. Kole (2016). The effects of asset forfeiture on policing: a panel approach. *Economic Inquiry* 54(1), 558–575.
- Knowles, J., N. Persico, and P. Todd (2001). Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy* 109(1), 203–229.
- Lahiri, P. and M. D. Larsen (2005). Regression analysis with linked data. *Journal of the American Statistical Association* 100(469), 222–230.
- Linn, E. (2009). *Arrest Decisions: What Works for the Officer?* Number 5. Peter Lang.
- Linos, E. (2018). More than public service: A field experiment on job advertisements and diversity in the police. *Journal of Public Administration Research and Theory* 28(1), 67–85.
- Lipsky, M. (1971). Street-level bureaucracy and the analysis of urban reform. *Urban Affairs Quarterly* 6(4), 391–409.
- Loeffler, C. E., J. Hyatt, and G. Ridgeway (2019). Measuring self-reported wrongful convictions among prisoners. *Journal of Quantitative Criminology* 35(2), 259–286.
- Lofstrom, M. and S. Raphael (2016). Prison downsizing and public safety: Evidence from california. *Criminology & Public Policy* 15, 349.
- Makowsky, M. D. and T. Stratmann (2009). Political economy at any speed: what determines traffic citations? *The American Economic Review* 99(1), 509–27.
- Makowsky, M. D. and T. Stratmann (2011). More tickets, fewer accidents: How cash-strapped towns make for safer roads. *The Journal of Law and Economics* 54(4), 863–888.
- Mas, A. (2006). Pay, reference points, and police performance. *The Quarterly Journal of Economics* 121(3), 783–821.
- Mas, A. and A. Pallais (2019). Labor supply and the value of non-work time: Experimental estimates from the field. *American Economic Review: Insights* 1(1), 111–26.
- Mastrobuoni, G. (2013). Police and clearance rates: Evidence from recurrent redeployments within a city. *Unpublished Working Paper*.
- Mastrofski, S. D. (2004). Controlling street-level police discretion. *The Annals of the American Academy of Political and Social Science* 593(1), 100–118.
- Merola, M. (1981). The bronx pre-arraignment experience. *Public Productivity Review*, 250–252.

- Mocan, N. (2008). What determines corruption? international evidence from microdata. *Economic Inquiry* 46(4), 493–510.
- Morris, C. N. (1983). Parametric empirical bayes inference: theory and applications. *Journal of the American Statistical Association* 78(381), 47–55.
- Moskos, P. (2008a). The better part of valor: court-overtime pay as the main determinant for discretionary police arrests. In *Law Enforcement Executive Forum*, Volume 8, pp. 77–94.
- Moskos, P. (2008b). *Cop in the hood: My year policing Baltimore’s eastern district*. Princeton University Press.
- Mukherjee, A. (2017). Impacts of private prison contracting on inmate time served and recidivism. *Available at SSRN 2523238*.
- Oettinger, G. S. (1999). An empirical analysis of the daily labor supply of stadium vendors. *Journal of Political Economy* 107(2), 360–392.
- Olken, B. A. and R. Pande (2012). Corruption in developing countries. *Annual Review of Economics* 4(1), 479–509.
- Olmstead, M. (2018). Nypd officers accused of “collars for dollars” arrests at end of shifts to rake in overtime pay.
- Ouss, A. (2015). Incentives structures and criminal justice. *Social Science Research Network Paper No. 2685952*.
- Owens, E. (2020). The economics of policing. *Handbook of Labor, Human Resources and Population Economics*, 1–30.
- Persico, N. and P. Todd (2006). Generalising the hit rates test for racial bias in law enforcement, with an application to vehicle searches in wichita. *The Economic Journal* 116(515), F351–F367.
- Potter, G. (2013). The history of policing in the united states, part 1.
- Prendergast, C. (2001). Selection and oversight in the public sector, with the los angeles police department as an example. Technical report, National Bureau of Economic Research.
- Prendergast, C. (2003). The limits of bureaucratic efficiency. *Journal of Political Economy* 111(5), 929–958.
- Prendergast, C. (2007). The motivation and bias of bureaucrats. *The American Economic Review* 97(1), 180–196.
- Punch, M. (2013). *Police Corruption: Exploring Police Deviance and Crime*. Routledge.
- Rehavi, M. M. and S. B. Starr (2014). Racial disparity in federal criminal sentences. *Journal of Political Economy* 122(6), 1320–1354.

- Skolnick, J. H. (2008). Policing should not be for profit. *Criminology & Public Policy* 7, 257.
- Smith, C. A. (2018). *Audit of the Dallas Police Department's Off-Duty Employment Program*.
- Smith, D. A. and C. A. Visher (1981). Street-level justice: Situational determinants of police arrest decisions. *Social Problems* 29(2), 167–177.
- Sun, I. Y. (2003). Officer proactivity: A comparison between police field training officers and non-field training officers. *Journal of Criminal Justice* 31(3), 265–277.
- Tahamont, S., Z. Jelveh, A. Chalfin, S. Yan, and B. Hansen (2020). Dude, where's my treatment effect? errors in administrative data linking and the destruction of statistical power in randomized experiments. *Journal of Quantitative Criminology*.
- Tonin, M. and M. Vlassopoulos (2015). Are public sector workers different? cross-european evidence from elderly workers and retirees. *IZA Journal of Labor Economics* 4(1), 11.
- Verma, A. (2016). A turning point in mass incarceration? local imprisonment trajectories and decarceration under californias realignment. *The Annals of the American Academy of Political and Social Science* 664(1), 108–135.
- Vila, B. (2006). Impact of long work hours on police officers and the communities they serve. *American Journal of Industrial Medicine* 49(11), 972–980.
- Vila, B., G. B. Morrison, and D. J. Kenney (2002). Improving shift schedule and work-hour policies and practices to increase police officer performance, health, and safety. *Police Quarterly* 5(1), 4–24.
- Walker, S. (1977). A critical history of police reform. *Lexington, MA: Lexington Books*.
- Weisburst, E. (2020). Whose help is on the way? the importance of individual police officers in law enforcement outcomes. *Working Paper*.
- Wilson, J. (1989). Bureaucracy: What government agencies do and why they do it.
- Worrall, J. L. and T. V. Kovandzic (2008). Is policing for profit? answers from asset forfeiture. *Criminology & Public Policy* 7(2), 219–244.
- Wu, X. and C. Lum (2017). Measuring the spatial and temporal patterns of police proactivity. *Journal of Quantitative Criminology* 33(4), 915–934.

Table 1: Summary Statistics for Officer-Years

	(1) Mean	(2) SD	(3) Min	(4) Max	(5) N
Number 911 Calls	616.22	430.31	1.00	2511.000	6145
Number Arrests	27.61	31.40	0.00	440.000	6145
Share Guilty Conviction	0.37	0.21	0.00	1.000	5661
Share Felony Arrests	0.21	0.18	0.00	1.000	5661
Share Civilian-Initiated Arrests	0.64	0.27	0.00	1.000	5661
Annual Salary	64522.928	12877.48	42940.8	113451.96	6098
Overtime Pay	5717.219	9004.43	0	127114.549	6145
Officer Female	0.16	0.37	0.00	1.000	6098
Officer Black	0.26	0.44	0.00	1.000	6098
Officer Hispanic	0.20	0.40	0.00	1.000	6098
Bachelor's Degree +	0.49	0.50	0.00	1.000	6098
Officer Tenure	11.55	9.30	0.00	46.000	6098

Note: Table presents summary statistics at the officer-year level.

Table 2: 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
Warrant	0.18	0.00	0.05	-4.88	0.40
Disorderly Conduct	0.15	0.00	0.07	-4.88	0.35
Assault	0.13	0.26	0.33	-4.84	0.47
Narcotics & Drugs	0.11	0.41	0.80	-4.66	0.47
Theft-Retail	0.06	0.11	0.75	-4.67	0.45
Public Intoxication	0.05	0.00	0.15	-4.54	0.30
Not Coded	0.04	0.00	0.18	-4.76	0.32
Trespass	0.04	0.01	0.83	-5.02	0.41
Dwi	0.03	0.12	0.83	-4.58	0.54
Traffic	0.02	0.00	0.55	-4.68	0.54

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. A shift-hour of -8 represents the first hour of an officer's shift; a shift-hour of -1 represents the final hour of an officer's shift.

Table 3: Summary Statistics, Off-Duty Employment

	(1)	(2)	(3)	(4)	(5)
<i>A. Officer-Year Level</i>	Mean	SD	Min	Max	N
Number of Off-Duty Shifts	45.27	63.03	0.00	540.000	7891
Average Off-Duty Shift Length	5.60	1.86	0.00	20.000	5618
<i>B. Daily Police-Shift Level</i>	Mean	SD	Min	Max	N
Any Off-Duty Shift					
After-Shift	0.098	0.30	0.00	1.00	607497
Before-Shift	0.065	0.25	0.00	1.00	607497
During-Shift	0.004	0.06	0.00	1.00	607497
Regular Off-Duty Shift					
After-Shift	0.069	0.25	0.00	1.00	607497
Before-Shift	0.050	0.22	0.00	1.00	607497
During-Shift	0.002	0.05	0.00	1.00	607497

Note: Table presents summary statistics for data on off-duty employment. Panel A presents information tabulated at the officer-year level — the number of off-duty work shifts and the length (in hours) of an off-duty work shift. Panel B presents information tabulated at the day-by-shift level. Here, we report the share of shifts that are either preceded or proceeded by any off-duty work shift or a regularly-scheduled off-duty work shift.

Table 4: Effect of Predicted Off-Duty Employment on Late vs. Early Shift Arrests

	Hourly Sample			Dispatch Sample			Daily Sample	
	(1) Any Arrest	(2) Non-Felony	(3) Felony	(4) Any Arrest	(5) Non-Felony	(6) Felony	(7) Overtime	
Regular Before × Early	-0.000785 (0.000403)	-0.000159 (0.000367)	0.000286 (0.000166)	0.000405 (0.000793)	0.000873 (0.000724)	-0.000531 (0.000349)		
Regular Before × Late	-0.0000825 (0.000421)	-0.000305 (0.000368)	0.000221 (0.000181)	-0.00220* (0.000908)	-0.00171* (0.000774)	-0.000477 (0.000392)		
Regular After × Early	-0.0000222 (0.000402)	-0.0000375 (0.000372)	0.0000321 (0.000141)	-0.000411 (0.000713)	-0.0000857 (0.000669)	-0.000358 (0.000282)		
Regular After × Late	-0.000715* (0.000340)	-0.000558 (0.000312)	-0.000138 (0.000124)	-0.00277*** (0.000803)	-0.00255*** (0.000716)	-0.000306 (0.000387)		
Regular Before Shift								-0.0336*** (0.00453)
Regular After Shift								-0.0543*** (0.00450)
Mean	0.017	0.014	0.003	0.032	0.026	0.006	0.283	
Observations	4895449	4895449	4895449	1571593	1571593	1571593	611911	

Standard errors in parentheses

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


Note: Table presents estimates from a series of regressions of an indicator variable for whether an officer made an arrest during a given shift on an indicator for whether the officer was predicted to work a regular off-duty shift either before or after his or her police shift. These regressions correspond with equation (3). Results are presented separately for early-shift arrests and late-shift arrests as well as for felony and non-felony arrests. The first three columns pertain to our hourly sample of all arrests. In columns (4)-(6), we report estimates for the sample of calls which were initiated by a citizen call for service. Standard errors are clustered at the officer and vision-by-month level.

Table 5: Calibrated Model Parameters

Parameters	Estimates	Description
ρ	0.037	Probability of guilt
μ	1.555	Mean of signal for guilty individuals (officer ability)
λ	0.154	Tradeoff for arresting guilty and not arresting innocent individuals
c_{ot}	-0.200	Intercept value/cost of working overtime
c_{od}	-0.007	Intercept value/cost of working off-duty
b	0.00011	Value of \$1 of overtime/off-duty pay
ϕ	0.791	Probability of processing an arrest
p	0.761	Per-period probability of continued arrest processing
p_{start}	0.347	Probability of beginning work on patrol
p_{ot}	0.284	Probability of receiving overtime when not making an arrest
Parameters	Estimates	Description
λ/b	1382.0	Dollar value of non-arrest of innocent person
$(1 - \lambda)/b$	7621.3	Value of arrest of guilty person
$E(v(a, s) a = 0)$	1345.5	Average value of non-arrest
$E(v(a, s) a = 1)$	2843.5	Average value of arrest
$E(\max_{a_t} v(a, s)) - v_{np}$	57.5	Average value of hour on patrol (relative to hour not on patrol)
$(\lambda/\phi - c_{ot})/b$	3547.6	Overtime pay needed for arrest of innocent suspect
$(1 - \lambda)/\phi b + (c_{ot} - c_{od})/b$ $+ Pay_{OT}$	8023.6	Off-duty pay needed for non-arrest of guilty suspect


Note: Table presents calibrated model parameters from a dynamic model of a police officer's decision to arrest a suspect.

Figure 1: The Probability and Amount of Overtime Pay by Shift-Hour of Arrest




Note: The left-hand panel plots the probability that an arrest made in a given hour of an officer's shift leads to overtime pay. The -8 hour corresponds with the first hour of the shift; the -1 hour corresponds with the final hour of the officer's shift. Statistics are computed separately for felony arrests and non-felony arrests. The dotted blue line represents the probability that overtime hours are worked for a shift in which no arrest is made. The right-hand panel plots the mean number of overtime hours worked for an arrest that is made in a given hour of an officer's shift. Again, statistics are presented separately for felony and non-felony arrests and the dotted blue line represents the mean number of overtime hours worked when a shift in which no arrest was made leads to an overtime spell. 95 percent confidence intervals are provided for each statistic.

Figure 2: Arrest Frequency and Court Conviction by Shift-Hour




Note: The left-hand panel plots the probability that an arrest is made in a given hour of an officer's shift. The -8 hour corresponds with the first hour of the shift; the 0 hours corresponds with the final hour of the officer's shift. We also plot the conditional probability that an arrest was made given that an arrest was made earlier in the officer's shift. The right-hand panel plots the probability that an arrest results in a criminal conviction for either a misdemeanor or a felony offense for an arrest made in each hour of an officer's shift. 95 percent confidence intervals are provided for each statistic.

Figure 3: Arrest Propensities By Hour and Shift




Note: Figure plots the average number of arrests per officer by hour of day, separately for officers in each of seven overlapping shifts. As is evident from the figure, in every hour of the day there are officers from at least two overlapping shifts on patrol.

Figure 4: 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 (2). 95 percent confidence intervals, computed using standard errors that are clustered at the officer and 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 5: Arrest Propensity From a 911 Call



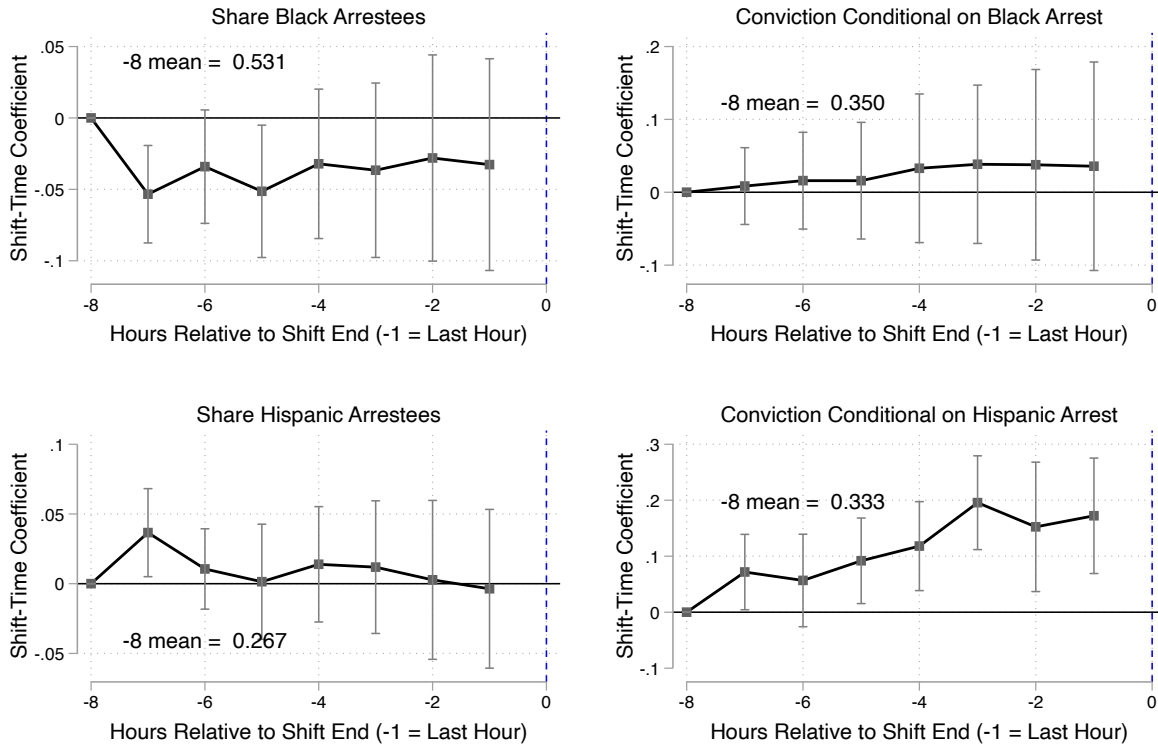
Note: Figure plots each shift-hour coefficient from a regression of a binary arrest indicator on a vector of shift-hour indicator variables, conditional on the fixed effects described in equation (2). The unit of analysis is the service call. 95 percent confidence intervals, computed using standard errors that are clustered at the officer and 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 6: Arrest Frequency and Court Conviction By Type



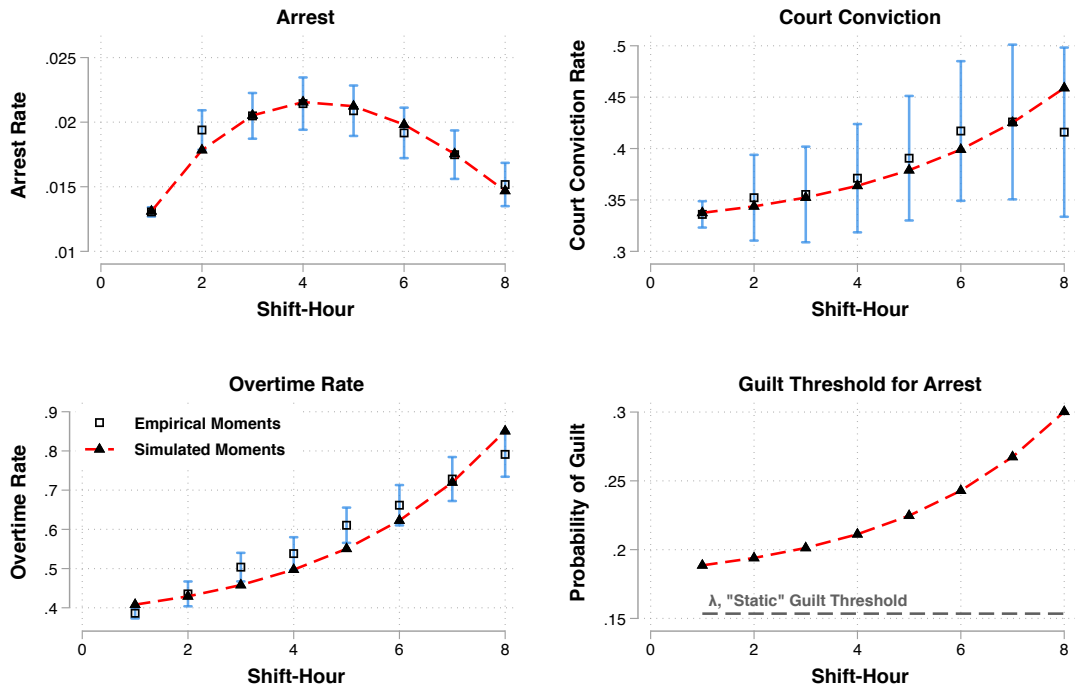
Note: Figure plots each shift-hour coefficient from a regression of a given outcome variable on a vector of shift-hour indicator variables, conditional on the fixed effects described in equation (2). We present estimates separately for arrests for violations, misdemeanor arrests, felony arrests and felony 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 officer and 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 7: Arrest Frequency and Court Conviction By Race




Note: 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 (2). 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, computed using standard errors that are clustered at the officer and 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 8A: Model Fit, Primary Estimates




Note: Figure plots empirical versus simulated moments as a test of model fit. Estimates are presented for arrest incidence, the court conviction rate and the rate with which an arrest leads to an overtime spell. We also plot the simulated guilt threshold as a function of shift-time.

Figure 8B: Model Fit, Off-Duty Estimates



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

Figure 9: Counterfactual Estimates of the Time-Path of the Arrest Rate and Guilt Threshold Used by Officers



Note: Figure plots the time-path of the rate of arrest (the left-hand panel) and the guilt threshold used by offices (the right-hand panel) throughout the officer's shift under three scenarios. The green line refers to a scenario in which officers are indifferent between working overtime and leaving work. The blue line refers to a scenario in which overtime pay is 6x the officer's wage rather than 1.5x the officer's wage. The red line is the estimated value from our model.

ONLINE APPENDIX

Appendix A: 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, 2020), 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:

$$Late_i = \alpha_{o(i)} + \phi_{dwh} + \theta_{dws} + \xi_{dm} + \epsilon_i \quad (6)$$

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 **Appendix Figure 4** which plots both the unadjusted and shrunk distributions of the officer fixed effects.³² We next use the shrunk 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

³²Because 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.

forms the basis for a claim that such behavior is motivated by the desire to secure access to overtime pay.

We explore these relationships in **Appendix Figure 5** 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 shrunk 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 2. 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.

Table A.1: “First Stage” Regression: Relationship Between Predicted and Actual Off-Duty Work

	(1)	(2)
	Before OD	After OD
Off-Duty Before Shift	0.570*** (0.00857)	0.0492*** (0.00564)
Off-Duty After Shift	0.00176 (0.00265)	0.578*** (0.00793)
Mean	0.065	0.099
Observations	611913	611913

Standard errors in parentheses

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

Note: Table presents estimates from a regression of an indicator for actual 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 officer and division-by-month level.

Table A.2: Balance Table: 911 Calls for Service by Shift-Hour

	(1)	(2)	(3)	(4)
	Log Calls	Log Serious Calls	Log Calls	Log Serious Calls
1hr From End	-0.00544 (0.00293)	-0.00233 (0.00234)		
2hr From End	-0.00923** (0.00301)	-0.00476* (0.00242)		
3hr From End	-0.00964*** (0.00272)	-0.00141 (0.00216)		
4hr From End	-0.0102*** (0.00252)	-0.000889 (0.00189)		
5hr From End	-0.00998*** (0.00268)	-0.00204 (0.00208)		
6hr From End	-0.00575** (0.00189)	-0.00353* (0.00177)		
7hr From End	-0.00205 (0.00186)	-0.00140 (0.00155)		
Before-Shift Off-Duty			-0.000630 (0.00128)	-0.000333 (0.00112)
During-Shift Off-Duty			-0.00742* (0.00352)	-0.00217 (0.00269)
After-Shift Off-Duty			0.00479** (0.00155)	-0.0000679 (0.00102)
Mean	2.048	1.455	2.048	1.455
Observations	4895449	4895449	4895449	4895449
F-value	2.960	1.498	4.849	0.227
F-test	0.005	0.166	0.003	0.878

Standard errors in parentheses

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

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 (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.3: Robustness of Estimates to Alternative Models, Primary Results

<i>F-tests of last three hours</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Arrest – Hourly Sample							
F-value	198.09	19.49	18.85	19.19	19.24	19.28	18.47
F-test	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Arrest – Dispatch Sample							
F-value	5.08	5.23	5.09	4.87	4.76	4.79	3.04
F-test	(0.0019)	(0.0016)	(0.0019)	(0.0025)	(0.0029)	(0.0028)	(0.0293)
Officer FE	X	X	X	X	X	X	X
Division FEs							
Division-Hour	X						
Division-Hour-DOW		X	X	X	X	X	
Division-Date						X	
Division-Date-Hour							X
Shift FEs							
Division-Shift			X				
Division-Shift-DOW				X	X	X	X
Time FEs							
Year-Month	X	X	X	X			
Year-Month-Division					X		
DOW	X						

Note: Table presents the F -statistic on the shift-hour coefficients for the final three hours of an officer's shift arising from a series of regressions of an arrest indicator on a vector of shift-hour indicator variables, conditional on various sets of fixed effects. For the hourly sample, the leave-out shift-hour is five hours from the end of shift, as we empirically observe this hour as the arrest-rate peak. The leave-out shift-hour for the dispatch sample is the first hour of the shift. Statistics are presented for the hourly sample as well as the dispatch sample. Each column corresponds to a different model. Standard errors are clustered at the officer and division-by-month level.

Table A.4: Robustness of Estimates to Alternative Models, Arrest “Quality” Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Guilty	0.0086 (0.0034)	0.0090 (0.0034)	0.0107 (0.0048)	0.0121 (0.0059)	0.0129 (0.0057)	0.0158 (0.0061)	0.0327 (0.0135)
Sentenced	0.0071 (0.0028)	0.0072 (0.0029)	0.0110 (0.0039)	0.0129 (0.0051)	0.0138 (0.0049)	0.0195 (0.0054)	0.0348 (0.0086)
Felony	0.0070 (0.0028)	0.0065 (0.0029)	0.0076 (0.0043)	0.0021 (0.0050)	0.0023 (0.0049)	0.0018 (0.0064)	0.0099 (0.0080)
Guilty — F Arrest	0.0068 (0.0055)	0.0079 (0.0066)	0.0025 (0.0056)	-0.0003 (0.0108)	0.0014 (0.0102)	-0.0358 (0.0213)	.
Guilty — M Arrest	0.0021 (0.0053)	0.0021 (0.0053)	0.0029 (0.0122)	0.0344 (0.0167)	0.0386 (0.0163)	0.0200 (0.0154)	.
Guilty — V Arrest	0.0048 (0.0025)	0.0041 (0.0026)	0.0057 (0.0041)	0.0167 (0.0071)	0.0172 (0.0070)	0.0245 (0.0077)	.
Officer FE	X	X	X	X	X	X	X
Division FEs							
Division-Hour	X						
Division-Hour-DOW		X	X	X	X	X	
Division-Date						X	
Division-Date-Hour							X
Shift FEs							
Division-Shift			X				
Division-Shift-DOW				X	X	X	X
Time FEs							
Year-Month	X	X	X	X			
Year-Month-Division					X		
DOW	X						

Note:

Table A.5: 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.6: Robustness of Off-Duty Estimates to Alternative Models

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	arrest	arrest	arrest	arrest	arrest	arrest	arrest
Regular Before × Early	-0.000148 (0.000395)	0.0000135 (0.000397)	0.0000107 (0.000397)	-0.0000305 (0.000400)	-0.0000785 (0.000403)	-0.0000209 (0.000398)	-0.000109 (0.000402)
Regular Before × Late	-0.000208 (0.000425)	-0.0000645 (0.000427)	-0.00000303 (0.000426)	-0.0000345 (0.000428)	-0.0000825 (0.000421)	-0.0000101 (0.000429)	-0.0000280 (0.000427)
Regular After × Early	-0.0000916 (0.000412)	-0.0000156 (0.000404)	-0.00000478 (0.000405)	-0.0000299 (0.000400)	-0.0000222 (0.000402)	-0.000103 (0.000407)	-0.000140 (0.000400)
Regular After × Late	-0.000755* (0.000348)	-0.000772* (0.000340)	-0.000715* (0.000338)	-0.000723* (0.000339)	-0.000715* (0.000340)	-0.000799* (0.000342)	-0.000683* (0.000335)
Mean	0.017	0.017	0.017	0.017	0.017	0.017	0.017
Observations	4895451	4895451	4895451	4895451	4895449	4895449	4893837
Officer FE	X	X	X	X	X	X	X
Division-Hour FE		X	X	X	X	X	X
Division-Hour-DOW FE		X	X	X	X	X	
Division-Date FE	X	X	X	X	X	X	X
Division-Date-Hour FE							X
Division-Shift FE			X				
Division-Shift-DOW FE			X	X	X	X	X
Year-Month FE	X	X	X	X			
Year-Month-Division FE					X		
DOW FE	X						

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 an indicator for whether a shift-hour is early (first four hours) or late (second four hours) in an officer's shift. Each model conditions on a different set of fixed effects.

Table A.7: Effect of Predicted Off-Duty Employment, Short versus Long Pre-Shift Work

	Hourly Sample			Dispatch Sample		
	(1) Any Arrest	(2) Non-Felony	(3) Felony	(4) Any Arrest	(5) Non-Felony	(6) Felony
Regular Before × Early × Short	0.000182 (0.000498)	0.000119 (0.000451)	0.00000313 (0.000210)	0.000199 (0.000982)	0.00108 (0.000914)	-0.000940* (0.000384)
Regular Before × Late × Short	0.0000205 (0.000463)	-0.000124 (0.000407)	0.000137 (0.000223)	-0.00266** (0.00102)	-0.00185* (0.000935)	-0.000767 (0.000457)
Regular Before × Early × Long	-0.000518 (0.000603)	-0.000625 (0.000564)	0.0000708 (0.000234)	0.000772 (0.00124)	0.000535 (0.00113)	0.000168 (0.000601)
Regular Before × Late × Long	-0.000265 (0.000709)	-0.000615 (0.000606)	0.000361 (0.000264)	-0.00142 (0.00144)	-0.00147 (0.00123)	0.0000280 (0.000595)
Regular After × Early × Short	0.0000336 (0.000403)	-0.0000574 (0.000363)	0.000107 (0.000174)	-0.000643 (0.000906)	-0.000454 (0.000824)	-0.000252 (0.000362)
Regular After × Late × Short	-0.000444 (0.000386)	-0.000369 (0.000345)	-0.0000456 (0.000153)	-0.00267** (0.000905)	-0.00244** (0.000820)	-0.000332 (0.000456)
Regular After × Early × Long	-0.000121 (0.000625)	-0.0000218 (0.000588)	-0.0000814 (0.000212)	-0.0000862 (0.000879)	0.000423 (0.000824)	-0.000498 (0.000399)
Regular After × Late × Long	-0.00112* (0.000469)	-0.000843 (0.000437)	-0.000276 (0.000170)	-0.00291* (0.00115)	-0.00270** (0.000979)	-0.000262 (0.000556)
Mean	0.017	0.014	0.003	0.032	0.026	0.006
Observations	4895449	4895449	4895449	1571593	1571593	1571593

Standard errors in parentheses

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

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). Each model conditions on a different set of fixed effects. Standard errors are clustered at the officer and division-by-month level.

Table A.8: Robustness of Off-Duty Estimates to Alternative Models, “Close” versus “Far” Off-Duty Work

	arrest	arrest	arrest	arrest	arrest	arrest	arrest
Regular Before	-0.000110	0.0000334	-0.0000172	-0.0000602	-0.000119	-0.0000856	0.000170
× Early × Close	(0.000539)	(0.000548)	(0.000557)	(0.000549)	(0.000552)	(0.000549)	(0.000552)
Regular Before	0.000194	0.000228	0.000280	0.000367	0.000308	0.000429	0.000332
× Late × Close	(0.000603)	(0.000612)	(0.000612)	(0.000618)	(0.000605)	(0.000611)	(0.000592)
Regular Before	-0.000223	-0.0000543	-0.0000151	-0.0000533	-0.0000882	-0.0000822	-0.000375
× Early × Far	(0.000526)	(0.000530)	(0.000523)	(0.000533)	(0.000533)	(0.000530)	(0.000534)
Regular Before	-0.000593	-0.000372	-0.000313	-0.000449	-0.000484	-0.000459	-0.000424
× Late × Far	(0.000525)	(0.000523)	(0.000521)	(0.000522)	(0.000519)	(0.000526)	(0.000543)
Regular After	0.0000172	0.0000544	0.0000824	0.000109	0.000128	0.0000632	-0.0000707
× Early × Close	(0.000446)	(0.000432)	(0.000431)	(0.000427)	(0.000425)	(0.000432)	(0.000415)
Regular After	-0.000877*	-0.000991*	-0.000976*	-0.000960*	-0.000941*	-0.00101*	-0.000881*
× Late × Close	(0.000389)	(0.000384)	(0.000381)	(0.000383)	(0.000387)	(0.000389)	(0.000379)
Regular After	-0.000321	-0.000165	-0.000202	-0.000346	-0.000363	-0.000471	-0.000258
× Early × Far	(0.000705)	(0.000698)	(0.000688)	(0.000686)	(0.000693)	(0.000697)	(0.000705)
Regular After	-0.000428	-0.000247	-0.0000980	-0.000136	-0.000153	-0.000273	-0.000195
× Late × Far	(0.000536)	(0.000528)	(0.000521)	(0.000511)	(0.000510)	(0.000518)	(0.000518)
Mean	0.017	0.017	0.017	0.017	0.017	0.017	0.017
Observations	4895451	4895451	4895451	4895451	4895449	4895449	4893837
Officer FE	X	X	X	X	X	X	X
Division-Hour FE		X	X	X	X	X	X
Division-Hour-DOW FE		X	X	X	X	X	X
Division-Date FE	X	X	X	X	X	X	X
Division-Date-Hour FE							X
Division-Shift FE			X				
Division-Shift-DOW FE			X	X	X	X	X
Year-Month FE	X	X	X	X			
Year-Month-Division FE					X		
DOW FE	X						

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 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.9: Effect of Predicted Off-Duty Employment on Selected Outcomes

	(1)	(2)	(3)	(4)
	Guilty	AnySentence	Felony	Mis/Fel
Regular Before Shift	0.00505 (0.0106)	0.00249 (0.00927)	0.00757 (0.00787)	-0.00460 (0.00900)
Regular After Shift	-0.00896 (0.00975)	-0.000893 (0.00871)	0.00345 (0.00763)	-0.00593 (0.00851)
Mean	0.355	0.247	0.187	0.724
Observations	62042	62042	62042	62042

Note: Table presents estimates from a series of regressions of a given outcome on an indicator for whether the officer was predicted to work a regular off-duty shift either before or after his or her police shift. Models condition on the fixed effects described in equation (1). Standard errors are clustered at the officer and division-by-month level.


Table A.10: Calibrated Model Parameters, Alternative Specifications

Parameters	Baseline	Court Errors	Wald Estimates	Both	Description
ρ	0.037	0.022	0.037	0.022	Probability of guilt
μ	1.555	1.943	1.554	1.943	Mean of signal for guilty individuals (officer ability)
λ	0.154	0.086	0.154	0.086	Tradeoff for arresting guilty and not arresting innocent individuals
c_{ot}	-0.200	-0.217	-0.212	-0.230	Intercept value/cost of working overtime
c_{od}	-0.007	-0.008	-0.013	-0.015	Intercept value/cost of working off-duty
b	0.00011	0.00012	0.00020	0.00022	Value of \$1 of overtime/off-duty pay
ϕ	0.791	0.792	0.791	0.792	Probability of processing an arrest
p	0.761	0.761	0.761	0.761	Per-period probability of continued arrest processing
p_{start}	0.347	0.337	0.347	0.337	Probability of beginning work on patrol
p_{ot}	0.284	0.284	0.284	0.284	Probability of receiving overtime when not making an arrest
Parameters					Description
λ/b	1382.0	702.8	761.4	384.7	Dollar value of non-arrest of innocent person
$(1 - \lambda)/b$	7621.3	7515.5	4193.2	4111.2	Value of arrest of guilty person
$E(v(a, s) a = 0)$	1345.5	693.8	741.2	379.7	Average value of non-arrest
$E(v(a, s) a = 1)$	2843.5	2569.4	1564.4	1405.5	Average value of arrest
$E(\max_{a_t} v(a, s)) - v_{np}$	57.5	61.9	31.6	33.9	Average value of hour on patrol (relative to hour not on patrol)


Note: Table presents calibrated model parameters from a dynamic model of a police officer's decision to arrest a suspect.

Figure A.1: Police, Arrests and Index Crimes: U.S. Cities, 2016

(A) Police Officers and Index Crimes




(B) Police Officers and Index Crime Arrests




Note: Figure plots the natural logarithm of index crimes (Panel A) and the natural logarithm of index crime arrests (Panel B) against the natural log of sworn police officers for cities with populations over 250,000 residents in 2016. Data on crimes, arrests and police manpower come from the Federal Bureau of Investigation's Uniform Crime Reports accessed from [Kaplan \(2019\)](#).

Figure A.2: 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 (2). 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 officer and division-by-month level, are provided for each statistic.

Figure A.3: Probability of Overtime Pay by Shift-Hour, Regression Adjusted



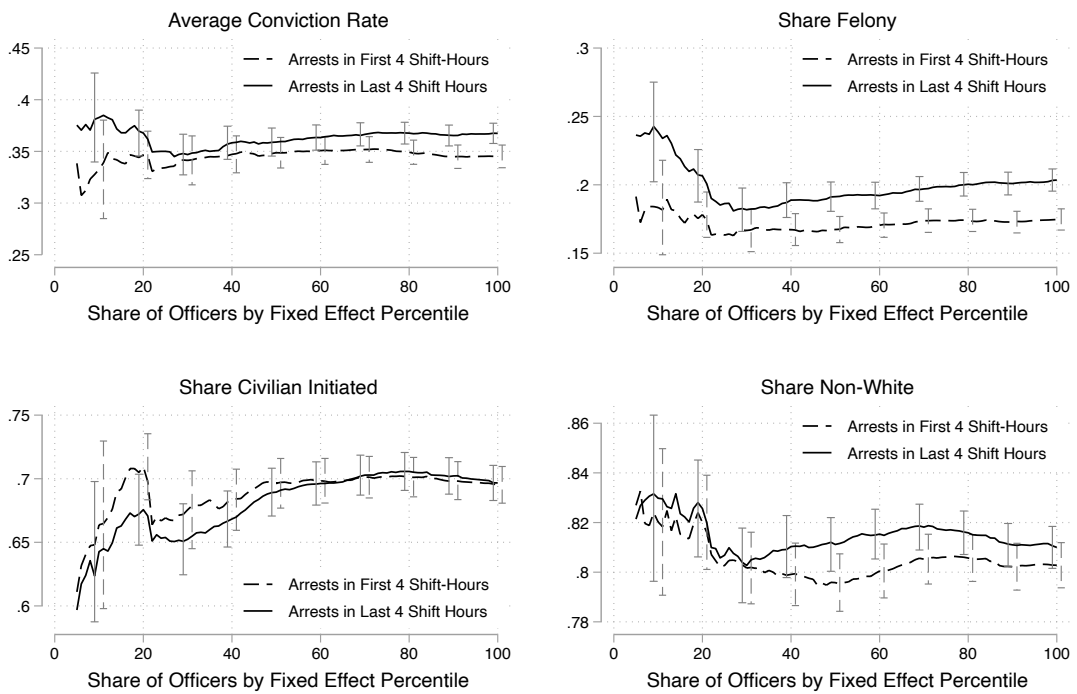
Note: The left-hand figure plots the number of arrest charges as a function of given hour of an officer's shift. The right-hand figure plots the probability of a court conviction in each hour of an officer's shift, separately for the full sample and for shift ending between 4:00pm and 10:00pm on Friday and Saturdays. All estimates condition upon the fixed effects in equation (2). 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 officer and division-by-month level, are provided for each statistic.

Figure A.4: Distribution of Officer Fixed Effects



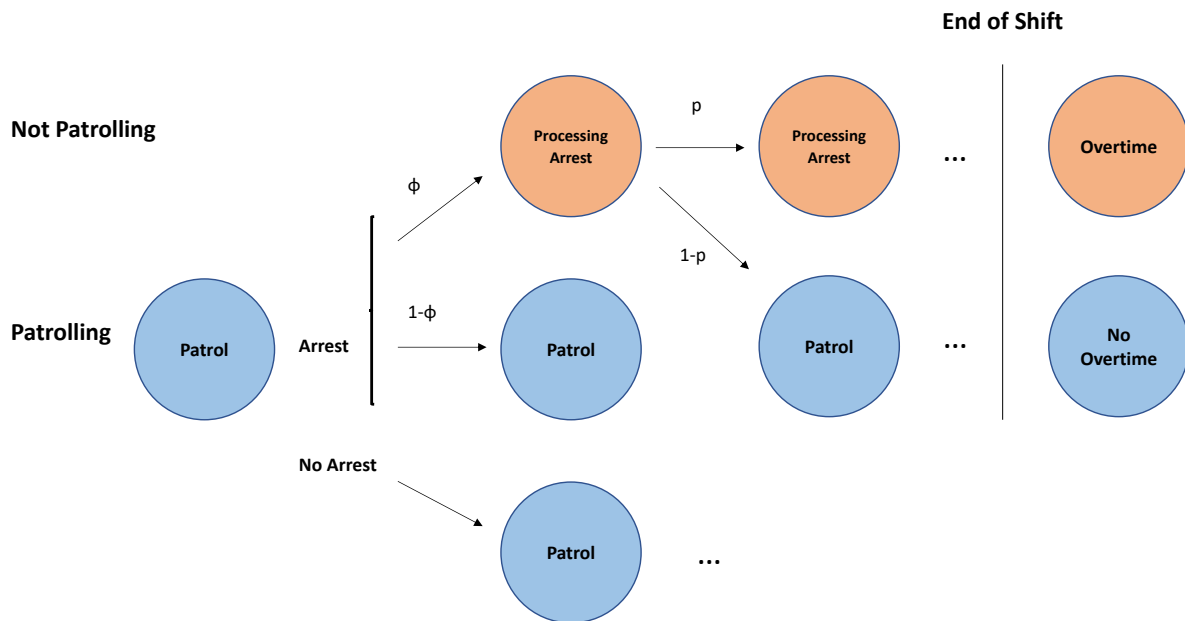
Note: Figure plots the the unadjusted distribution (the solid line) and the shrunk distribution (the dotted line) of officer fixed effects with respect to late-shift arrest activity. The fixed effects are estimated using a regression of whether an arrest occurs in the final two hours of an officer's 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.5: 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.6: Visualization of Model



Note: Figure presents a schematic of the dynamic model of an officer's arrest decision presented in Section 6.

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.