

A Political-Economic Model of Police Administration

Tom S. Clark, University of Chicago

A primary challenge for the administration of public safety is the implementation of policy that facilitates effective policing and satisfies political pressure over outcomes. I develop a model of police administration that captures the tension between policing outcomes and policing tactics. The model evaluates how police officers' incentives to violate policy limits on their behavior complicate delegation by leaders. Police leaders are motivated to provide discretion so officers may use appropriate tactics for whatever situation they face but also motivated to ensure that police do not abuse discretion and comply with limits on their behavior. I use the model to study the effects of political control of policing policy and insulation of police officers from accountability.

In the summer of 2014, police in Ferguson, Missouri, shot and killed a civilian named Michael Brown. Brown's death served as a flash point to catalyze public concern about the use of force by police against those they protect. The event brought about sustained public discontent and energized scrutiny of the relationship between law enforcement agencies and people of color in the United States. Indeed, a lawsuit brought by the Department of Justice against Ferguson resulted in a consent decree under which Ferguson agreed to reform its use of force policy (among other things). The Department of Justice was concerned about the use-of-force policy because it believed Ferguson gave officers too much leeway to use force in situations when it might not be necessary and did not sufficiently encourage officers to de-escalate confrontational policing events. At the same time, no criminal charges were brought against Darren Wilson, who shot Brown, and he was not fired for the incident.¹

Events such as those that took place in Ferguson raise serious questions about how American policing operates. How do policing policies affect the way police carry out their jobs? I build a model of police administration that focuses on the delegation problem between police leaders and police officers and captures some of the unique aspects of the police hierarchy.

The model I develop isolates a key tension at the heart of police administration and provides analytic insight into a problem of political delegation that is not unique to law enforcement settings but perhaps is at its highest stakes in that

context. Policy leaders make policy, a key component of which is the limits on officers' tactical options—what behaviors are permissible and which are not. In doing so, they balance potentially competing incentives. On one hand, they are incentivized to produce high-quality policing outcomes, which necessarily entails endowing officers with sufficient discretion to respond as necessary to policing events. For example, in the event of a dangerous person wielding a knife and threatening bystanders, officers need to have sufficient tools to disarm and subdue the suspect. On the other hand, police leaders face strong incentives to avoid excessive police abuse, which entails imposing limits on what police may permissibly do. In the wake of George Floyd's murder, departments around the country revisited their policies with respect to various kinds of tactics. However, these leaders make policy choices in the shadow of external, political pressure to insulate officers from discipline. The model examines how a police leader balances those competing incentives in a world where institutional constraints, such as police unions or policy makers, impose limits on the leader's ability to discipline officers or set the officers' level of discretion.

The core logic underlying the analysis is that police officers may face strong incentives to violate policy, because of their own bias, extreme policing events that require extreme tactics, or institutional protections that insulate the officer from accountability. For a police leader allocating discretion to officers, those incentives complicate policy making. The police

Tom S. Clark (tomsclark@uchicago.edu) is the David and Mary Winton Green Professor of Political Science, University of Chicago.

An online appendix with supplementary material is available at <https://doi.org/10.1086/729933>.

1. Wilson did later resign from the Ferguson Police Department.

leader must construct a policy that properly balances the officer's incentives to comply with policy against the incentives to ignore limitations on his behavior. The analysis reveals that the effects of many policy prescriptions, such as strict limits on officers' tactical options, can have nonobvious, unanticipated effects. Designing optimal policing policy cannot be done without considering the sets of biases that the officers themselves bring to their work. The result of the interaction of these various forces is a complex mix between the nature of police officers, the political environment in which policy is set, and the kinds of policing in which an offer engages.

THE POLITICS OF POLICING POLICY

Police executives, among whom police chiefs are the prototypical example, face a number of pressures that shape how they develop policing policy and, in turn, what kind of policing we observe take place. First, they face internal pressures arising from the complex nature of police work. Second, they face external pressures that arise from the larger institutional and political environment in which they develop and implement policy. Third, police executives face hierarchical pressures from politicians to whom they are accountable and the officers and potential officers they employ. Those forces, in turn, affect the resources, support, and options available to police leaders.

Police and police work

The literature on policing widely recognizes that officers are quintessential "street-level bureaucrats" (Lipsky 1971). However, police officers differ from more commonly studied bureaucrats in two aspects. First, they are near-omnipresent representatives of the government; a police officer is perhaps the most likely government representative with whom a citizen will come into contact. Second, while officers typically carry out important, but ultimately bureaucratic, routine work that is heavily governed by policy, they differ from other bureaucrats in that there is always a chance that any given policing event could turn violent or otherwise dangerous.

That police serve primarily routine bureaucratic functions but do occasionally confront high-stakes, violent events complicates how administrative police governing bureaucratic discretion is designed. A wide body of literature has studied the strategic (and nonstrategic) calculus of delegation in a bureaucracy (for overviews, see Gailmard and Patty 2012; Huber and McCarty 2004). Focusing on how a political principal can constrain an administrative bureaucrat charged with developing or implementing policy, those models contemplate both agency loss due to informational asymmetries (e.g., Bendor and Meiorowitz 2004; Epstein and O'Halloran 1994; Huber and Shipan 2006; McCubbins, Noll, and Weingast 1987) and efforts to incentivize bureaucrats to either develop or apply

their expertise to their administrative functions (e.g., Alesina and Tabellini 2007; Gailmard and Patty 2007; Huber and McCarty 2004; Maskin and Tirole 2004). The core idea throughout this literature is that bureaucrats play some role in administering or implementing policy, and that role can influence outcomes. The challenge becomes, then, ensuring outcomes with which the principal is satisfied.

Less often do models focus on the specifics of agent efforts—the ways in which they exercise effort. Usually, that is not a limiting consideration because the tasks in which a bureaucrat engages are relatively routine, and their implications for outcomes are straightforward and well captured by the models of administration and implementation.² However, in the context of policing, the particular tactics officers employ and actions they undertake while implementing policy are in fact often as salient and politically relevant as the outcomes they produce. (Officer discretion is exceptionally complex and touches on myriad aspects of their job, including where, when, and how to police [e.g., Mastrofski 2004].) It is not sufficient, then, to simply apply canonical models of agency and oversight to study how street-level bureaucrat police carry out their duties, especially in the more dangerous and violent settings.

Among the aspects of discretion they oversee, police leaders decide on what kinds of tactics police officers may use. In one salient example, in the aftermath of George Floyd's death, national attention turned to the use of choke holds, and police departments around the country adopted, or considered adopting, policies that prohibited officers from using them. Indeed, deciding on the degree of tactical discretion officers have at their disposal is a central element of policing policy. Use-of-force policies are an extremely prominent component of police governance and training, and high-profile instances of police abuse frequently involve the misuse of tactical tools or weapons.

Of course, police chiefs care not just about compliance with policy but also the quality of policing outcomes that are produced, and I turn to that below. However, it is important to underscore that policing tactics matter independently of policing outcomes for a variety of reasons. First, what tactics police employ can be related to the way the public perceives officers and the legitimacy conferred to the police (e.g., Hawdon, Ryan, and Griffin 2003). Second, policing tactics are likely to be a source of variation in safety—both for officers and for civilians (e.g., Meyer 1992). And, importantly, there is reason to believe that policing policy does in fact shape the kinds of

2. A notable exception to this general description is the literature on whistle-blowing, which explicitly considers actions that bureaucrats might undertake that are not directly about implementation but can clearly affect outcomes (see, e.g., Ting 2008).

tactics police officers use when they are performing police work (Mummolo 2018b). Indeed, in contrast with most other street-level bureaucrats, the particular tactics and choices officers make in the line of duty is a core component of oversight and accountability. The public, politicians, and police leaders care about not just the provision of public safety but also the specific means by which it is provided. Thus, inducing compliance with use-of-force policy is a primary objective for police leaders and is a core element of how they are evaluated by their own superiors.

What is more, it is not just police leaders for whom policing outcomes and policing tactics are important. Politicians, to whom police leaders are accountable, have very strong reasons to prioritize both policing outcomes and policing tactics. Public safety is a signal issue for politicians, especially local politicians. For this reason, politicians have a strong interest in the quality of policing outcomes produced by the police. In addition, they have an interest in the ways in which police produce their outcomes (i.e., the tactics police use). For example, excessive police behavior and tactics can lead to public outcry and protests, and evidence suggests such protests can affect local politicians' electoral fortunes (e.g., Gillion 2020).

Police unions and policing policy

Of course, police leaders are not alone in making either policy or personnel decisions. A salient feature of policing—both in the world of public safety policy makers and in public debates about policing—is the power of police unions. These institutions exert considerable influence on both how police leaders shape policy and how they exercise oversight of officers on the job.

One of the primary ways through which police unions shape policy is by controlling the nature and process of officer discipline. Mostly, their influence is indirect. Unions seek to affect the broader legal context in which officers operate, which creates incentives for police leadership when it designs the scope and condition of officer work. At least early on, “union contracts [had] little to say about how police officers actually do their job. This is in part because police unions, like other public and private sector unions in the 1960s, focused on gaining the power to bargain over economic issues, promotions, and discipline” (Fisk and Richardson 2017, 739). More important, in addition to negotiating the terms of labor contracts directly, police unions influence policing policy by becoming involved in electoral politics (e.g., Fisk and Richardson 2017). One way in which unions become particularly involved is through lobbying efforts related to police-reform legislation. For example, police unions have been very successful in preventing so-called sunshine laws that would make public police officers' disciplinary records (e.g., Bies 2017). They have also been particularly successful in opposing civilian

oversight boards (e.g., Bies 2017; Keenan and Walker 2004; Sklansky 2008), although in recent years police are increasingly accepting some role for civilian oversight. Unions not only blocked police reforms that would have increased police accountability but were also successful in advancing legislation that shielded police from accountability and often had the effect of protecting them against discipline for misconduct (e.g., Bies 2017). Among the most notable of such efforts has been the adoption of police officer “Bill of Rights” in many states.

Politics and police administration

Police administrators confront the complex nature of policing policy and the pressure of police unions in the context of a complex political economy. In particular, police administrators face potentially competing pressures from political principals to whom they are accountable and economic constraints imposed by the labor market for police officers (e.g., Matusiak 2016).

Police chiefs are under constant pressure from their political principals to produce high-level public safety, which creates very significant constraints on their ability to administer policing policy (e.g., Geller 1985). The primary manifestation of this pressure is an incentive to produce high-quality policing outcomes. (This can be thought of as different from policy compliance, which is more concerned with the particular tactics officers use.) A policing outcome refers to the qualitative and quantitative features of the resolution of an event in which the police are involved. A wide range of academic and policy literature has focused on how best to measure police performance and public perceptions of police performance (e.g., Davis et al. 2015; Schafer, Huebner, and Bynum 2003). One lesson from that literature is that police work is complex and involves many potentially orthogonal dimensions. However, the literature also suggests that we can imagine a dimension of police performance that corresponds roughly to public safety. For example, a policing event might be an incident in which a driver is suspected to be under the influence of alcohol. The outcome includes things such as whether a citation was issued, whether an arrest was made, whether the officer was assaulted, whether force was used against the civilian, and so forth. For simplicity, we can imagine that any given community—including the public, police leaders, and police officers—can rank all possible policing outcomes from better to worse. At one end are nonconfrontational, peaceful interactions between civilians and officers in which something approaching an idealized version of “justice” is done in the name of advancing public safety. At the other end are policing outcomes that involve, for example, unjustified force, severe injustice, or extreme confrontation and do not obviously promote public safety.

One consequence of political oversight and the complex environment in which police executives work is that police chiefs tend to have very short tenures, leading to a virtual “revolving door” for police leadership (Balfe 2015; Rainguet and Dodge 2001). Given the preferences politicians have about policing policy and outcomes, it is not surprising that police chiefs should find that their job security is weak. The principals to whom they are accountable have little interest in retaining police executives and strong incentives to take steps that at least signal to their constituents that they are paying attention to issues of public safety.

While police executives face pressure from above, they also face pressure from below. Ideally, a police chief would identify the characteristics needed for an officer to carry out optimal policing policy and go hire such officers. In reality, the police chief is constrained by the ability to hire qualified police officers. In recent years, the media has paid increasing attention to the “vacancy crisis” in American policing. The job of being a police officer is an inherently dangerous one, and, whether or not it is true, there is a perception that policing has become even more dangerous in the twenty-first century. As in other sectors, the competitiveness of a particular job opportunity waxes and wanes with broader economic conditions. Indeed, the labor market constraints on police departments can be severe, with excessive shortages of officers from Alaska to Connecticut (US Department of Justice 2018). The reasons for the drop remain relatively unstudied, but journalists and pundits have been quick to point to Ferguson, high-profile instances of fatal ambushes of officers, and low unemployment as sources of declining numbers of officers (e.g., Dewan 2017). Though, police departments have also long decried the dearth of high-quality candidates, and police academies typically have very high “wash-out” rates (e.g., Bernstein 2019; Hicks 1986). As a result, police chiefs try to build complex policing policy, subject to interventions from powerful actors, such as police unions, and oversight from politicians demanding high-level performance. However, the availability of potential police offers can limit their ability to achieve these goals.

A MODEL OF POLICE ADMINISTRATION

The model I develop is focused on the choices of a police executive, represented by a police chief, and a police officer. While unions, politicians, and the labor market play important roles in how those actors behave, I isolate these two key actors for analytic clarity.

Players and sequence of play

The game is played between a police chief, C , who represents the police command staff, and a police officer, O . The police chief is responsible for the administration of policing

policy, whereas the police officer carries out police work. The sequence of play is as follows.

1. Nature selects a political environment, characterized by a tuple, $\langle \alpha, \beta, \gamma \rangle$, an officer type, $\theta \in \mathbb{R}$, and a wage for police officers, $b > 0$. These choices are publicly observable.
2. The police chief selects a policy, π , which is a pair (τ, ρ) , where $\tau \geq 0$ is the degree of tactical discretion police officers have, and $\rho \geq 0$ is the degree of protection officers have from disciplinary actions arising from their use of those tactics.
3. A policing event, $e \in \mathbb{R}_+$ is realized.
4. The police officer chooses a tactic, $t \geq 0$. I refer to higher t 's as “more aggressive” tactics.
5. A policing outcome, ω , is realized, which is a function of t , e , and a random shock, ε . (ε captures the unpredictable and unverifiable details that transpire in the midst of a policing activity and is centered in 0.)
6. The police chief observes ω , t , and e and decides whether to retain the officer. Let $r \in \{0, 1\}$ represent the choice to retain the officer, where $r = 1$ denotes the decision to retain.
7. Payoffs are realized, and the game ends.

Policing outcomes

Policing outcomes are determined by a combination of the policing event, e , and the tactics the police officer chooses, t . For simplicity, I assume policing outcomes are given by

$$\omega = (e - t)^2 + \varepsilon \quad (1)$$

so that “higher” police outcomes correspond to increasing mismatch between the policing event and the tactics chosen by the officer.

Policing policy

The first component of policing policy, τ , is a threshold in the tactical choice space for the officer. Substantively, this corresponds to a distinction between tactical choices that are allowed and those that are never allowed. For example, recent debates in the United States have considered whether police officers should be allowed to use choke holds on suspects they are detaining. In the model, changing policy to allow or disallow choke holds would correspond to raising or lowering τ , respectively. This representation is admittedly stark in that it assumes tactical discretion is simply a single threshold that applies to all policing events, τ . In a richer model, the choice of policing policy might be the choice of a function that maps specific events into specific tactical ranges. However, it is

unclear how enforceable event-specific policies can be insofar as police officers themselves observe the policing events and often have latitude in describing what has transpired and the justification for their tactical choices. By focusing on a single limit on discretion, the model focuses analytic attention on the kinds of policing policy debates that typically dominate political debate, such as whether some tactics are ever permissible.

Beliefs and information

There are two sources of incomplete information. First, both the police chief and the officer are uncertain about the random shock that influences policing outcomes, ε . They share a common belief that ε is a random variable that is uniformly distributed on the interval $[-\zeta, \zeta]$. Second, before making policing policy, the police chief is uncertain about what policing event the officer will face but believes the distribution of policing events is drawn from a distribution, $g(\cdot)$, over the positive real number line, \mathbb{R}_+ . I assume that the distribution of policing events is exogenous and independent of policing policy. Of course, it is possible that policing policy affects the kinds of policing events that arise, through multiple mechanisms, which I discuss below. In analyzing the model, I assume that $g(\cdot)$ is the exponential distribution, characterized by a parameter, $\lambda > 0$, but it is sufficient that the distribution have a sufficiently long right tail (i.e., that there are rare but extreme policing events) for all of the results to follow.

Preferences and utility

The police chief has preferences over (a) policing outcomes, (b) officer compliance with policy, and (c) officer discipline and vacancies. The external political-economic environment shapes how much she cares about policing outcomes and political oversight. Political pressure on the police chief determines the incentives she has to produce high-quality policing outcomes and induce compliance with policy. Politicians—especially leaders in cities and municipalities—often place a lot of weight on ensuring high-quality public safety. By contrast, the power of police unions shapes the incentive for officer protection from discipline and the consequences of vacancies on the force. The cost of a vacancy might be complex in nature. For example, during the protests in June 2020, several officers around the country were terminated for misconduct. One officer in Buffalo was fired after pushing an elderly man to the ground, causing him to be hospitalized. As a result, the entire team of officers serving in the crowd control unit resigned from that unit. Garrett Rolfe in Atlanta was terminated and charged with murder after he fatally shot Rayshard Brooks. Nearly five-sixths of the Atlanta police department walked out in response. These kinds of actions impose extreme costs on

police leadership, as they struggle to maintain important police functions. However, there are also more small-scale costs, such as the effort that has to be allocated to recruiting and training new officers. And, of course, there are political repercussions that come with having an understaffed force. I assume the chief's utility function is given as follows:

$$U_c(\tau, \rho, r; \theta) = -r \cdot (\alpha \cdot \omega + \beta \cdot \mathbf{1}_{t > \tau} (t - \tau)^2 + \rho^2 - \gamma \cdot \rho) - (1 - r) \cdot c. \quad (2)$$

The parameter α captures the degree of interest the police chief has in the policing outcomes, whereas the parameter γ measures the degree of pressure on the police chief to protect officers from termination. This interest could be grounded in the challenges of the local labor market or the police chief's external interactions, such as with politicians or labor unions. The parameter β , by contrast, measures the police chief's interest in compliance with policy limitations on the officer's tactical discretion. The $\mathbf{1}_{t > \tau}$ function is an indicator that takes on the value 1 if $t > \tau$ and 0 otherwise. This term in the utility function captures the notion that police chiefs do not benefit from compliance with policy but only incur costs from violations of policy. It is important to note that this term is distinct from the term capturing policing outcomes, which is shaped by how well tailored the officer's behavior is to the situation he confronts.

The police officer has preferences over the tactics he uses and his career. Specifically, the officer wants to use policing tactics that are well matched to the policing event he encounters, but I allow the officer's type to influence his aversion to either under- or overmatching the event. That is, some officers might be more willing to risk an overresponse, whereas other officers might be more willing to risk an underresponse. One interpretation of the officer's bias is as a metric of the officer's preference for "overpolicing" or "underpolicing," although popular discourse about those concepts often involves more complex notions of policing practice. I assume that the police officer's preference over policing outcomes and compliance with policy is indirect and flows through his incentive to retain his job. That is, the officer's direct preference is over how he responds to policing outcomes, given his bias, not the actual policing outcome, ω , or compliance with policy, r :

$$U_o(t; \pi, \omega, \theta) = -(e + \theta - t)^2 + r \cdot b. \quad (3)$$

Table 1 summarizes the notation used in the model.

EQUILIBRIUM ANALYSIS

To analyze the game, I characterize a subgame perfect Nash equilibrium. Define a *setting* as $\Gamma = \langle \alpha, \beta, \gamma, \theta, c, b, \zeta, g \rangle$. Now, consider first the police chief's decision whether to retain

Table 1. Summary of Notation Used in the Model

Parameter	Definition	Choice	Definition
θ	Officer type (bias)	τ	Chief's choice of tactical limits
b	Wage (benefit) of police job	ρ	Chief's choice of protection from termination
e	Policing event	π	The tuple (τ, ρ)
α	Weight police chief places on outcomes	t	Officer's choice of tactic
β	Weight police chief places on compliance with policy	r	Chief's choice whether to retain officer
γ	Pressure to protect police officers		
ε	Random shock in policing outcomes		
2ζ	Support for ε		

the officer. Given π and ω , the police chief will retain the officer if and only if $\omega \leq c - \beta \cdot \mathbf{1}_{\{t > \tau\}}(t - \tau)^2 - \gamma(\rho^2 - \rho)/\alpha$.

That is, by construction, the friction of the labor market creates room for an officer to choose tactics that result in outcomes that he would be fired for were there no employment friction. This feature of the model is by design; important policing reform debates often identify the protection police unions, for example, provide to officers whose behavior warrants termination.³

Remark 1. The policy chief cannot credibly commit to firing officers who violate limits on their tactical discretion.

Given this decision rule for the chief, an officer who confronts a policing event optimally selects a tactic. Because the officer observes e but not ε , at the time he chooses a tactic, he does so subject to some uncertainty about the ultimate outcome. Therefore, he forms an expectation about the likelihood he will be fired for a chosen tactic. Specifically, the officer weighs the risk that the random component of the policing event is large enough to justify his termination. Given a tactic, t , the police chief will be willing to retain the officer whenever

$$\varepsilon \leq \frac{c - \beta \mathbf{1}_{\{t > \tau\}}(t - \tau)^2 + \gamma \cdot \rho - \rho^2}{\alpha} + 2te - t^2 - e^2. \quad (4)$$

There exists a unique subgame perfect Nash equilibrium to the model. In equilibrium, policing policy induces a partitioning of the policing event space into three regions. At the far right—for the most extreme policing events—the police officer finds policy too constraining, relative to the benefits of

retaining his job as a police officer. For events in this region, the officer simply chooses his ideal tactic for each event. At the far left—for the least extreme policing events—the police officer does not find policing policy, τ , constraining and chooses a tactic that balances his ideal tactical choice against the incentive to produce high-quality policing outcomes. This tactical choice, then, is partially debiased by his career incentives. In the middle range, the officer chooses a tactic that violates policy (i.e., $t > \tau$ in equilibrium). However, as compared to less extreme policing events, the police officer must further debias his tactic in order to account for not just the quality of policing outcomes produced by the tactic but also the fact that his choice violates policy. The choice of τ balances the policing outcomes given each of these three kinds of tactics. The choice of ρ simply creates friction in the disciplinary process in response to external pressure (γ). As I show below, the pressure to protect police officers therefore has only second-order effects on the degree of tactical discretion the police chief adopts.

Proposition 1. In the unique subgame perfect Nash equilibrium to the game, the police chief chooses a policy in which the degree of protection depends only on the pressure to protect officers, γ , but the degree of tactical discretion is a function of the distribution of policing events, g ; the officer type, θ ; and the political context, α and β . The officer's tactical choice is weakly increasing in the event extremity. The officer is probabilistically retained in equilibrium.

Proof. All proofs are gathered in the appendix.

To develop intuition for the police officer's equilibrium strategy, it will be convenient to distinguish between two kinds of officers. The first is an officer with type $\theta > 0$, and the second is an officer with type $\theta < 0$. I refer to these two kinds of officers as “high-type” and “low-type” officers, respectively. In

3. For example, during the protests that took place after the killing of George Floyd, Oates (2020) described the challenges police chiefs face in terminating police officers who had engaged in misconduct.

the unique subgame perfect equilibrium to the model, the officer's behavior depends on his type, with qualitative differences between officers who are high type and those who are low type. High-type officers prefer to use tactics that are biased toward more aggressive responses to policing events than are appropriate. Low-type officers prefer to use tactics that are biased toward less aggressive tactics than are appropriate. As we will see, policing policy and career concerns have opposing effects on how officers respond to policing events.

Notice first that if an officer places no weight on his career concerns, then his strategy is trivial—he matches his tactical choice to his perception of the policing event. He simply selects a tactic to match his perception of the policing event (i.e., he chooses $t = e + \theta$). That tactic can be an overresponse or an underresponse, depending on the officer's bias, θ . In fact, as we will see, there are conditions under which an officer ignores his career incentives. However, it is possible that the incentives created by the police chief's oversight induce him to choose something other than his sincere preferred tactic. In particular, one concern might be that for his preferred tactic, a policing outcome might be so bad that it exceeds the degree of protection he receives from policy, ρ . That concern will determine how far he deviates from his sincere preference. There is, though, a second force that influences the officer's decision. Once the officer's strategic (or sincere) tactical choice exceeds the limit on his discretion, τ , then the police chief experiences disutility, exacerbating the constraining effect of career concerns for the officer. Therefore, how far the officer deviates from his preferred tactical choice can depend also on whether his strategic tactical choice violates policing policy.

By the same logic, it is possible that the degree to which an officer must deviate from his preferred tactic is too large relative to his interest in retaining his job. Simply put, there can arise policing events that are so extreme that for a high-type officer, the degree to which he must deviate downward in order to make the chief willing to retain him—to stay close enough to policy limits on his discretion—becomes intolerable, and the officer decides to forgo his career incentives in order to respond to a policing event as he sees fit. Alternatively, there can also be policing events sufficiently extreme that in order to produce high enough quality policing outcomes, a low-type officer might have to adopt a tactic that he is unwilling to choose. That is, given a low θ , there can be policing events that require a sufficiently aggressive tactic that the officer would prefer to ignore his career concerns and choose his preferred tactic. In this way, both types of officers can face policing events that are too extreme for them to be constrained by oversight. High-type officers decide to ignore limits on their tactical discretion, whereas low-type officers decide to ignore limitations on the effectiveness of their policing tactics.

The qualitative effect of these incentives depends on whether the officer is a high type or a low type. The officer's equilibrium strategy is a function of the event, e , given by

$$t(e)^* = \begin{cases} \frac{2\zeta(e + \theta) + eb}{b + 2\zeta} & \text{if } e \leq \underline{e}(\theta) \\ \frac{b(\alpha e + \beta\tau) + 2\alpha\zeta(\theta + e)}{b(\alpha + \beta) + 2\alpha\zeta} & \text{if } \underline{e}(\theta) < e \leq \bar{e}(\theta) \\ e + \theta & \text{else.} \end{cases} \tag{5}$$

(Throughout the analysis, I focus on the case in which $\underline{e}(\theta) > 0$. However, it is possible that $\underline{e}(\theta) < 0$. In that case, the equilibrium space simplifies so there is only one partitioning, $\bar{e}(\theta)$, which is always positive, but equilibrium behavior does not change.)

Consider first the strategy for sufficiently low policing events (i.e., $e \leq \underline{e}(\theta)$). This tactic differs from the unconstrained tactic, $e + \theta$, by a margin of $\theta - [2\zeta(b + \theta + e)/b]$. So, depending on the benefit of retaining his job, a high-type officer chooses a less aggressive tactic than he would absent career incentives, whereas a low-type officer chooses a more aggressive tactic than he would absent career incentives. Nevertheless, the derivative of this tactic, with respect to e , is $b + 2\zeta/b > 0$, meaning the officer's tactic becomes more aggressive as the policing event becomes more extreme.

Consider next the strategy for more extreme policing events (i.e., $e > \underline{e}(\theta)$). For these events, the officer continues to moderate his tactical choice away from his sincere preference. However, the effect of career concerns is more pronounced. The marginal change in his policing tactic is greater than the marginal change in policing events, meaning that the officer's deviation from his sincerely preferred tactic grows as he faces more and more extreme policing events.

That effect gives rise to the third qualitative region of the policing event space—the region where the officer is no longer constrained by career incentives. As noted above, it is possible that the officer finds the necessary degree of deviation from his sincere preference too great, relative to the benefit of keeping his job. In this region, which exists for the most extreme policing events, the extent to which the officer must modify his preferred behavior to stay in compliance with policy is not warranted.

Figure 1 illustrates equilibrium tactics for an example high-type and low-type officer. In figure 1A, we see that for policing events that are extreme enough, high-type officers moderate their behavior, using less extreme tactics than they would absent policing policy constraints. For sufficiently extreme policing events, career incentives cannot constrain officer behavior. In figure 1B, we see that for policing events that are extreme enough, low-type officers can be induced to use more aggressive behavior, by pressure to produce high-quality




Figure 1. Representation of optimal policing tactic as a function of policing events, e : A, $\theta > 0$; B, $\theta < 0$. Solid gray diagonal line shows a tactic perfectly matched to the event. Solid black line shows equilibrium tactics.

policing outcomes. As with high-type officers, but for a different reason, for sufficiently extreme policing events, career incentives cannot induce the low-type officer to be aggressive enough to handle the event with which he is confronted. In brief, policing policy can induce increased consistency in police officers' tactical choices, but not for the most extreme events, in either direction.

To develop intuition for equilibrium policing policy, notice that the choice of τ does not pose direct costs to the police chief; rather, the effect of τ is through its effects on officer behavior and the cost of oversight caused by violations of policy (i.e., instances in which $t^* > \tau$). In addition, notice that

$$\frac{\partial \bar{e}(\theta)}{\partial \tau} = \frac{\partial \underline{e}(\theta)}{\partial \tau} = 1,$$

which means that the width of the interval over which a policing tactic is a function of τ is constant. However, what τ does affect is the *location* of the interval, $(\underline{e}(\theta), \bar{e}(\theta))$. Therefore, the police chief's utility is affected by the relative density of policing events between $\bar{e}(\theta)$ and $\underline{e}(\theta)$ that is determined by her choice of τ . Thus, the police chief's challenge is to choose a policing policy that optimizes her expected utility across these three regions.

In the first region, for $e < \underline{e}(\theta)$, the officer chooses a tactic, t^* , that is independent of τ . As the police chief decreases τ , $\underline{e}(\theta)$ decreases. The result is a marginal improvement in policing outcomes for each policing event that shifts from being below $\underline{e}(\theta)$ to above $\underline{e}(\theta)$. However, whereas the officer never violates policy, τ , for $e < \underline{e}(\theta)$, by definition tactics violate policy for events above $\underline{e}(\theta)$. Therefore, while decreasing τ leads to an improvement in utility from policing outcomes, it imposes a cost for events that a decrease in τ shifts to being above $\underline{e}(\theta)$.

Thus, one tension faced by the police chief concerns the marginal improvement in utility from inducing better outcomes at $\underline{e}(\theta)$ versus the marginal decrease in utility from inducing tactical violations of policy at $\underline{e}(\theta)$. This calculus will turn on the density of policing events at $\underline{e}(\theta)$ and the relative importance of policing outcomes as opposed to violations of policy (i.e., α vs. β). That tension is at the core of her choice of policing policy.

Of course, as τ changes, not only does $\underline{e}(\theta)$ change, so too does $\bar{e}(\theta)$. Increasing τ improves the police chief's utility for events at $\bar{e}(\theta)$, because the officer uses a tactic that more closely matches the event itself, as career incentives temper the effect of the officer bias below $\bar{e}(\theta)$. In addition, if the police officer is a high type, increasing τ improves the chief's utility, because it creates a marginal decrease in violations of policy by the officer. Recall, as shown in figure 1, high-type officers choose a more aggressive tactic for events over $\bar{e}(\theta)$ than for events just below $\bar{e}(\theta)$. However, if the officer is a low type, then he will choose a less aggressive tactic above $\bar{e}(\theta)$, and as a result increasing τ creates a marginal decrease in the police chief's utility for low-type officers. The optimal τ is the one for which these two marginal effects on the chief's utility offset each other exactly. There exists a unique policing policy that satisfies this condition, and so there is a unique τ^* for any given setting. (The uniqueness of τ^* depends on the assumption that g has sufficient skew. If extreme policing events are too common relative to less extreme policing events, it is possible that for some parameter values multiple equilibria can exist.)

While generic closed-form solutions for the equilibrium τ are not possible, I can calculate τ^* by assuming particular parameter values, which allows me to evaluate how equilibrium policing policy shifts with key parameters. Figure 2 shows




Figure 2. Equilibrium tactical discretion, τ^* , as a function of officer bias (θ) and pressure to protect officers (γ). A, Manipulates the value of holding a policing job. B, Manipulates the relative weights of policing outcomes and compliance with policy for the police chief. Plots assume particular values for model parameters not being varied in the plots.

the equilibrium τ for several different scenarios. Figure 2B compares τ^* when $\alpha > \beta$ against a world in which $\alpha < \beta$. The comparison, therefore, captures how the relative pressure to produce high-quality policing outcomes compared to compliance with policy affects optimal tactical discretion. On the X-axis of this plot is the degree of officer bias, θ . As the figure shows, when the police chief places more weight on high-quality policing outcomes than on compliance with policy, she gives more discretion when an officer is a high type than when the officer is a low type. This asymmetry is driven by different incentives each type of officer faces to violate policy, as discussed above. By contrast, as the chief cares more about compliance with policy, τ becomes more sensitive to the officer's absolute bias, rather than the direction of his bias.

By contrast, figure 2A shows τ^* as a function of the pressure to protect officers, γ . The solid line shows equilibrium policy given a relatively very high value of the job for the police officer (b), whereas the dashed line shows equilibrium policy given a relatively low value of the job for the police officer. Generically, the police chief chooses weakly more discretion for officers as the value of the job decreases; in order to induce officers to be constrained by career incentives, the police chief has to give the officers more discretion. Perhaps more interestingly, as pressure to protect officers increases, there is a nonmonotonic effect on equilibrium policing policy. Initially, protection for officers increases equilibrium discretion, but after a certain point, pushing for more protection induces a decrease in tactical discretion. This effect occurs because initial career protection for officers induces them to respond to events with appropriate tactics, and so increases in discretion do not necessarily lead to bad policing outcomes. However, with enough career protection, the effects of officer bias outweigh the performance benefits of tactical discretion and produce worse policing outcomes and less compliance with policy. Later, when I turn to discussing the relationship between police unions and policing outcomes, this effect will prove consequential.

The challenge for the police chief in choosing an optimal level of protection, ρ , is less complex. Recall that ρ merely creates friction in the oversight process, insulating officers against the random elements of policing events and allowing officers' own biases to influence their tactical choices. However, because the officers can control their tactical choices, they are able to choose a tactic that, in expectation, keeps the police chief indifferent about retaining the officer in the event that an adverse outcome occurs. Therefore, the police chief's choice of a level of protection only matters for the chief in situations when she might otherwise prefer to fire the officer. As a consequence, the optimal level of protection is simply the level of protection at which an otherwise unconstrained police chief would choose to set ρ . That choice is determined only by the relative utility associated with ρ , which is captured simply by the parameter γ .

It is notable that even though the choice of ρ is straightforward, its effects on equilibrium behavior are more complex, as it in turn influences τ and, consequentially, t . That is, as the equilibrium level of protection for the officer from termination increases, there is a nonmonotonic effect on the chief's choice of tactical discretion (fig. 2). Because increasing ρ makes it more difficult to terminate an officer, at low levels of γ , increasing γ creates incentives for officers to provide high-quality policing outcomes, because they anticipate being insulated in the event of an adverse random outcome or minor violation of policy. However, once γ reaches a certain point, further increasing tactical discretion merely creates incentives for officers to act on their own biases rather than respond to accountability incentives. As a consequence, the optimal τ decreases, in order to induce less biased tactical choices for those policing events for which an officer does respond to policy, offsetting the adverse effects of career protection for the more extreme policing events.

Proposition 2. In equilibrium, pressure to protect police officers (γ) has a nonmonotonic effect on officer

discretion. When pressure to insulate officers is low, increasing pressure to protect officers induces a greater degree of tactical discretion. Once officers are sufficiently insulated, further increasing pressure to protect officers decreases their tactical discretion, τ .

Proposition 2 illustrates an endogenously arising trade-off between career protection and discretion for police officers. Many models and studies of delegation capture similar trade-offs, often focusing on how principals can incentivize effective policy making or implementation by a motivated bureaucrat (e.g., Carpenter 2020; Gailmard and Patty 2007; Stephenson 2007). Here, the underlying mechanism is distinct. The trade-off emerges because the principal faces exogenous pressure to insulate the agent from accountability, which results in the principal holding the reins a bit tighter in the settings in which the agent is responsive to oversight. What is particular to the police administration setting is that external lobbying focuses so primarily on disciplinary issues that it puts pressure on the police chief to adopt career protection.

The equilibrium yields a number of comparative statics that can provide the basis for an analysis of the sources of policing disparities. In particular, the model helps understand how features of the officers and local policing demands affect equilibrium behavior. I now turn to these comparative statics before illustrating their implications for disparities in policing.

The officer, the job, and policing policy

The first set of comparative statics that shed light on the politics of police administration concerns the labor market for police officers. In particular, as we saw above, two important features that are central to police administration are the nature of officer bias and the desirability of the work. The former is measured by θ , whereas the latter is measured by b . I consider each in turn.

As we saw above, policy is a function of the police officer's bias. As an officer's type increases—as an officer prefers to overrespond rather than underrespond—equilibrium policing policy shifts. In particular, τ^* increases nonmonotonically in θ . Substantively, as an officer becomes more biased toward aggressive tactics, the police chief selects a higher threshold, τ^* . The chief is driven to do so by the incentive to induce the officer to stay within policy. If the threshold is too low, the officer will simply violate policy; if the officer does not have to moderate too much from his preferred tactic, he can be induced to comply with policy. However, once an officer becomes too biased, the optimal threshold decreases. The chief's incentive to produce good policing outcomes overcomes the incentive to induce compliance with policy for officers who are

too extremely biased. Somewhat counterintuitively, we do not see that discretion is maximized for unbiased officers. Rather, discretion is maximized for officers who are somewhat, but not too, biased toward aggressive tactics.

Lemma 1. Tactical discretion changes nonmonotonically in officer bias. As an officer becomes more biased in favor of aggressive tactics, the chief provides greater tactical discretion. Once an officer is sufficiently biased in favor of aggressive tactics, the chief's optimal level of discretion decreases.

At the same time, the officer's bias not only affects equilibrium policy but the qualitative nature of the equilibrium, because θ determines where the cut points $\underline{e}(\theta)$ and $\bar{e}(\theta)$ are located. The range of policing events for which the officer strategically moderates his tactical choice is moving to the left and decreasing in size as θ increases. Thus, the incentive to allow more overresponsively biased officers more discretion is more than offset by the officer's incentive to disregard career concerns and violate policy. The effect of increasing tactical discretion is simply to slow down the rate at which the officer approaches the point at which he is willing to forgo the benefit of retaining his job.

Figure 3 overlays the equilibrium tactic for a high-type and a low-type officer, holding constant policing policy. As the figure shows, high-type officers are more likely to ignore their career incentives and adopt their sincerely preferred tactic. At the same time, low-type officers are less likely to modify their equilibrium tactics because of constraints imposed by limits on their discretion (τ). However, what does become a binding constraint on low-type officers is the risk of producing sufficiently bad outcomes that the police chief prefers to terminate the officer. Thus, whereas high-type officers are more likely to be influenced by limits on tactical discretion, low-type officers are more likely to be influenced by expectations that they produce high-quality police work.

Further, as noted in proposition 1, it is possible that for a sufficiently small θ then $\bar{e}(\theta) < \underline{e}(\theta)$. If that is the case, then the equilibrium space simplifies, and there is simply a range, $e < \bar{e}(\theta)$, in which the officer chooses a tactic that is an underresponse mitigated by policing policy and policing concerns, and a range $e > \bar{e}(\theta)$, in which the officer chooses his ideal tactic for a policing event, $t = e + \theta$.

Consider next the effect of the desirability of being a police officer, measured by b . The attractiveness of the police officer's job creates an incentive for the officer to choose less biased tactics in the range for which he is constrained and also expands the range of policing events for which he is constrained. Therefore, the effects of b are qualitatively similar to




Figure 3. Comparison of equilibrium tactics for high- and low-type officers

the effects of officer bias. Increasing b leads to ceteris paribus less biased tactical choices by the officer.

Lemma 2. An officer chooses weakly less biased tactics as the benefit to retaining his job increases.

The consequence of this result is that policing policy will vary with the availability of outside employment options for officers. As I discussed above, a significant concern for policy makers and public safety officials in recent years has been the declining pool of candidates for officer positions, the high rate of turnover and “burn out,” and the consequent high vacancy rate in municipal police departments. Examples of departments facing shortages of officers who employ underqualified officers abound. In addition, programs such as the Department of Justice’s Vets to COPS initiative systematically channel some individuals (in this case, veterans) into local police forces. What is more, significant political and social forces have created pressure on departments to diversify their ranks to address discrepancies between the demographics of police forces and the communities they serve, but the efficacy of those efforts has been the subject of considerable debate.

The policing environment

Finally, it is important to note that a considerable source of substantive variation that affects the politics of police admin-

istration is the nature of the policing events to which officers must respond. There is notable variation not just in crime rates but in the kinds of crime that take place across jurisdictions. There is one straightforward comparative static that corroborates intuition. Denote $\tau(g)^*$ the equilibrium τ given the distribution of policing events, g . If g has first-order stochastic dominance over g' , then $\tau(g)^* > \tau(g')^*$. Loosely, if the distribution of policing events moves toward more extreme policing events, then equilibrium policing policy permits more tactical discretion.

Lemma 3. Equilibrium policing policy changes as the distribution of policing events changes. As more extreme policing events become more likely, the police chief gives the police officer more tactical discretion.

At the same time, as we have seen, equilibrium protection for officers, ρ^* , is just a linear function of the amount of pressure the police chief faces to protect officers, γ . As a result, there is no direct effect of the nature of the policing events an officer faces on the protection the police chief gives to him. Rather, policing events shape the kinds of tools an officer has at his disposal. Therefore, the degree of protection an officer receives from termination is best understood as a policy choice that rests on the presumption that the police chief is otherwise enabling the police officer to do his job effectively.

EXPLAINING DIVERGENT POLICE CULTURES

The model and results of the equilibrium analysis provide the foundations for a pair of thought experiments. We can consider two kinds of structural factors that give rise to policing disparities. First, the model allows us to examine how local political conditions affect policing policy and, by extension, policing behavior. I propose a thought experiment that imagines how local political conditions might influence the choice of policing policy that police leaders make. For example, we might imagine that increased attention to tactical excess by officers increases the pressure to contain officer tactics, forcing the police chief to choose a suboptimal constraint on officer discretion. Alternatively, we might expect that high-profile scandals, such as some of the salient examples of police killing civilians, increase pressure on police leaders to produce high-quality outcomes. Second, the model allows us to examine how one aspect of the labor market for police officers affects policing policy. Specifically, we focus on the role that police unions play in influencing both policy and personnel decisions and show the ways in which the local labor market for police can affect policing behavior and outcomes. In terms of the model, police unions might affect the parameter γ , which influences equilibrium protection for officers against termination.

Politics and policing behavior

Political circumstances surrounding policing can affect the model dynamics in two ways. First, they can alter key parameters, such as the weight that the police leader places on the quality of policing outcomes or the weight she places on officer compliance with policy. Second, political circumstances can impose exogenous constraints on the police leader's behavior, perhaps most notably her control over police tactics. Indeed, scholars know that policy changes might come about because of the political and social context in which the police chief develops policies and tactics (e.g., Bass 2000; McCall 2019) or via judicial decisions that induce change (e.g., Mummolo 2018b). Those kinds of events, in turn, often follow salient events that focus mass attention on policing.

An example of this dynamic is the “defund the police” movement that came into political salience in the summer of 2020, at least partially in response to the murder of George Floyd in May 2020. While the proposal is not always fully articulated or consistently stated, the core of this idea is that if the police continue to use their authority and resources to harm, rather than protect, civilians, then civilians should provide fewer resources and support to them—both fiscally and politically.

Many canonical models of electoral accountability would suggest that as popular support for a particular government policy, activity, or expenditure wanes, elected officials will be incentivized to decrease their support for the policy, activity, or

expenditure, even if the relationship between public opinion and policy making is imperfect (Ashworth 2012). In this setting, the implication of these arguments is that when the public becomes less supportive of the police, then politicians will have an incentive to rein in the police. That translates into increased pressure on police leaders to produce high-quality policing outcomes to keep officers in compliance with policy, as well as increased political constraints on the tools police officers have to use.

Consider an exogenous constraint on officers' tactical discretion such that they have fewer tools at their disposal than the police chief would choose in equilibrium (i.e., $\tau' < \tau^*$). As τ' decreases, there are countervailing effects on the kinds of police behavior we observe. Notice that as τ' decreases, the range of policing events for which the police officer's tactical choice is constrained by policy shifts toward less extreme policing events. The effects of exogenously decreasing τ away from equilibrium policy, therefore, depend on the officer's underlying bias. The result is less aggressive tactical choices for the marginal policing event for which the officer is constrained under equilibrium policy if the officer is biased toward under-responding but more aggressive tactical choices for the same event if the officer is biased toward over-responding. Thus, the effect of political efforts to constrain officer tactics depends on the qualitative nature of officer bias and can, counterintuitively, result in officers choosing either more or less aggressive tactics than they would choose under equilibrium policy.

Figure 4 illustrates the tactical choice for a high-type officer given equilibrium policing policy, τ^* , and policing policy given an exogenous shift downward in officer tactical discretion, to τ' . (The case for a low-type officer is simply the mirror image.) The 45° line shows $t = e$ (i.e., completely unbiased tactical choices, resulting in ideal policing outcomes for all policing events), and the gray curve shows an example distribution of policing events $g(e)$.

What we see here is that by reducing τ below its equilibrium level, there is a debiasing of police tactics for relatively less extreme policing events, combined with an exacerbation of bias in police tactics for the more extreme policing events. The total impact on policing outcomes, therefore, depends on both the magnitude of officer bias and the distribution of policing events. As more serious or extreme policing events become more common, decreasing officer tactical discretion results in worse policing outcomes. However, for more moderate or routine policing events, decreasing officer tactical discretion yields an improvement in the quality of outcomes. Thus, the effect of an exogenous reduction in officer discretion is one that improves what might be thought of as routine policing but potentially undermines the quality of police responses to extreme policing events. Precisely how big each of these changes




Figure 4. Police officer tactical choices given equilibrium policy, t^* , and exogenously lowered tactical discretion, τ'

is depends on the broader political environment as captured by model parameters α , β , and γ , as well as the distribution of policing events, g .

Implication 1. Political efforts to restrain police officer discretion lead to improvements in policing outcomes for more routine policing events but undermine the quality of policing in more extreme policing events. The magnitude of these effects depends on the broader political environment in which the police chief makes policy.

Thus, the model provides one theoretical foundation for variation in policing outcomes—the political environment in which police leaders make policy. Exogenous constraints on their authority over choices such as use-of-force policy can cut both ways. They can improve policing for some kinds of policing events but undermine it for others. The total effect on policing in a given community depends on the interaction of both the kinds of policing events officers confront as well as the broader climate in which a police chief and police officers operate. This result has implications for both policy reformers and empirical scholars interested in the effects of institutional changes to policing policy.

Unions, resources, and oversight

A second prominent subject in public and political discourse regarding policing disparities is police unions. Given the lit-

erature on the role police unions play, police unions can be understood as the force driving the incentive to protect officers from termination. They are strong advocates for greater protection for police officers—that is, they seek to increase the equilibrium ρ . (They can also directly affect the cost of terminating an officer [e.g., Oates 2020], but substantively the consequential and salient effects occur through their more systemic efforts to shape the contours of the police employment environment.) These effects occur both through the pressure police unions place on the political system, by lobbying for legislation that protects officers, as well as through ex post protection unions afford officers, through arbitration, litigation, and internal processes that can induce less severe punishment for officer misbehavior. An additional effect, although one practitioners often dismiss, is through unions' effects on working conditions and wages for officers, which are captured in the model by the parameter b .

As we saw in figure 2 and the discussion of proposition 2, increasing the police chief's incentive to protect officers from termination can induce a trade-off wherein the police chief chooses a policy that allows less tactical discretion in response to the behavioral incentives created for officer tactical choices by enhanced protection from termination. Some protection for officers creates incentives for officers to provide effective policing and also comply with policy. However, given sufficient protection, further increasing officer protection from termination exacerbates the effects of officer bias on policing

outcomes. The consequence is that the police chief will, in equilibrium, further limit the range of behavior that is acceptable under policy as she experiences further pressure to protect officers. (Recall that does not mean violating policy leads to termination, as we saw remark 1). Related, *ceteris paribus*, as γ increases, there is an ambiguous effect on equilibrium tactical choices. Notice that as γ increases, $\bar{e}(\theta)$ decreases. The officer chooses a weakly more biased tactic for all $e > \bar{e}(\theta)$ than he chooses for $e = \bar{e}(\theta)$. Therefore, an increase in γ creates a marginal increase in how biased police tactics are away from optimal tactics at the point where officers switch from constrained to unconstrained tactical choices (i.e., $|t^* - e|$ is increasing in γ at $\bar{e}(\theta)$). However, for those policing events below $\bar{e}(\theta)$, the effect of γ on the bias in policing events depends on the sign of θ . If $\theta > 0$, then increasing γ induces the officer to choose a more biased tactic, whereas if $\theta < 0$, then increasing γ induces the officer to choose a less biased tactic. Whether that positive effect below $\bar{e}(\theta)$ offsets the negative effect above $\bar{e}(\theta)$ depends on both the density of cases in the range $(\underline{e}(\theta), \bar{e}(\theta))$ and the location of $\bar{e}(\theta)$ induced by the other model parameters. As $\bar{e}(\theta)$ decreases, *ceteris paribus* it is more likely that increasing γ unambiguously leads to worse policing outcomes. The consequence of these various effects is that union pressure to insulate officers from termination unambiguously leads to worse policing outcomes if officers are sufficiently biased toward overresponding but has an ambiguous effect if officers are biased toward underresponding. Thus, an empirical implication of the model is that we should observe an inverse relationship between disciplinary insulation for officers and the amount of tactical discretion in policing policy. However, the effect of pressure to insulate officers on the quality of policing outcomes is more complicated and potentially ambiguous.

Implication 2. External pressure on the police chief to protect officers from termination reduces equilibrium permissible tactical discretion. If officers are biased toward overresponding, reducing tactical discretion unambiguously leads to worse policing outcomes. If officers are biased toward underresponding, the reducing tactical discretion has ambiguous effects on policing outcomes.

Whereas the effect of union efforts to insulate officers from termination are ambiguous, union efforts to improve working conditions and wages are more straightforward. As we saw in lemma 2, increasing the attractiveness of a police officer's job leads to unambiguously less biased tactical choices by the officer. Therefore, as police unions improve the benefit of holding a police officer's job—independent of protection from termination—officers produce better policing outcomes.

These two different features of union benefits create different mechanisms. Wages increase the value of holding a job and therefore unambiguously create an incentive to comply with policy. Yet, protection from discipline can create more room for officer bias to shape behavior and potentially lead to worse outcomes. Thus, an empirical implication of the model is that different forms of union lobbying will affect policing outcomes differently.

Implication 3. Disparities in policing outcomes can arise from union lobbying for protection from termination and improved wages and working conditions. Improved wages unambiguously increase the quality of policing outcomes, whereas protection from termination decreases policing outcomes under most conditions.

Of course, decreasing bias means that officers who are biased toward overresponding will use less aggressive tactics as the attractiveness of their job increases. However, officers biased toward underresponding will use more aggressive tactics as the attractiveness of their job increases. Thus, the qualitative effect on the tactical choices we observe is asymmetric for officer types but unambiguously increases the quality of policing outcomes.

Interestingly, these results are consistent with empirical evidence that unionized police officers generate more citizen complaints but that those complaints are less likely to be sustained upon investigation. One report compiled by the Bureau of Justice Statistics shows that while the rate of civilian complaints against police officers is higher in departments with collective bargaining, those complaints are sustained by the relevant bodies responsible for reviewing complaints in unionized departments at roughly half the rate as they were in nonunionized departments (Hickman 2006, 6). Those data are consistent with the theoretical prediction that union strength should lead to both worse policing outcomes and less disciplinary action against officers. Of course, many possible models make that prediction, and so this evidence serves merely to corroborate one empirical pattern implied by the theoretical account of policing disparities that follows from the model. However, the model's particular prediction that union protection will lead to more biased tactical choices by officers is one that can be subjected to empirical scrutiny with appropriate data. Unfortunately, data on both union strength and police tactical choices remain sparse and rarely systematic.

Stepping even further outside of the model, we might also imagine that an additional effect of union strength goes beyond career protection and the conditions of employment. One might imagine that union strength itself affects the qualitative characteristics of the individuals who are candidates in

the police labor market. If union strength affects θ —that is, if the kind of benefits that the police union offers shapes the characteristics of the workforce itself—then there can be a recursive, or self-reinforcing, dynamic whereby the dimensions of union strength determine the welfare implications of union strength. A union that is most successful at enhancing working conditions and wages (i.e., b) is more likely to have an unambiguously positive effect on policing outcomes than is a union that is successful at securing protection of officers from termination. Of course, as discussed above, unions are most active and successful at lobbying for police officer protection and are less involved in the working conditions for officers (e.g., Feuille and Juris 1976). I have not modeled the effects of unions on the labor market directly, but intuition and extant labor market theory suggest that career protection might be more important for less qualified officers, whereas wages and working conditions might be more important for more highly qualified officers. In terms of the model, that logic implies that unions are likely to lead to more biased officers, which, in light of the preceding analysis, suggests *ceteris paribus* worse policing outcomes. Of course, this remains conjecture, because the equilibrium implications of union effects on officer characteristics are beyond the scope of the model.

Conjecture 1. If union protection for officers results in more biased officers, then increased union strength leads to worse policing outcomes.

Finally, it is worth considering how we can extrapolate from the static model of a police chief and a single officer to policy making with an entire police force. It might seem that a police chief should simply target the officer modeled here—treating him as a representative officer of the force. However, it is important to note that the chief's utility is asymmetric in a particular way. Noncompliance with policy is most consequentially caused by high-type officers (those who prefer to overpolice). Underpolicing by low types does not symmetrically offset the disutility caused by high-type officers employing overly aggressive tactics because there is no proportional bonus from compliance. Alternatively, a police chief might simply target the most biased officer. However, that approach might induce a degree of tactical discretion that is too low for a substantial portion of the police force, especially if the highest type officer is very biased. (Recall, tactical discretion is nonmonotonically related to officer bias [lemma 1] and that suboptimal tactical discretion can have complex effects on policing outcomes [implication 1].) Thus, building policy for an entire police force will require consideration of the police chief's sensitivity to noncompliance (β) as well as calculations about the distribution of officers in order to calculate the aggregate con-

sequences of any given policy as applied to an entire police force.

DISCUSSION AND CONCLUSION

The effectiveness of modern policing practices is a subject of intense public and political debate and in recent years has come to occupy the interest of political scientists as well. Of particular interest are high-profile instances of seeming police misconduct accompanied by the use of tools and tactics that seem, at least to observers, excessive. Indeed, the prominence of civilian deaths at the hands of police in the United States has raised questions about the effect of police violence on the health of American democracy (e.g., Cohen et al. 2019; Gunderson et al. 2021; Lowande 2020; Mummolo 2018a) and has highlighted a long-standing concern about racial bias in law enforcement (e.g., Clark et al. 2023; Gelman, Fagan, and Kiss 2007; Knox, Lowe, and Mummolo 2020).

These sources of public and scholarly concern highlight the complexities of making policing policy and the disparities in policing observed across jurisdictions in the United States. The model in this article focuses on one particular complexity—the interaction of competing incentives facing policy makers with external political constraints. Powerful interests often push to insulate officers from discipline, and that insulation undermines police administrators' ability to hold officers accountable. As a result, policy makers must balance potentially competing incentives. On one hand, police leaders have a strong incentive to produce high-quality policing outcomes, which entails ensuring police officers, who respond to emergencies and make decisions on their own in real time, use appropriate tools and tactics for those events. However, those officers may have preferences for either over- or under-responding to any given event. On the other hand, police leaders have an incentive to craft policies that those (potentially biased) officers will comply with, to ensure lawful and consistent policing practice within the police force. An officer might choose to take advantage of the police leader's lack of a credible commitment to terminate him for violating policy, disregarding tactics he perceives as too constraining, potentially resulting in worse policing outcomes as well as flagrant violation of policy. Further, these police administrators navigate these competing concerns in the face of political pressure to insulate officers from discipline. Analysis of the model reveals one logic for how police leaders may navigate these competing concerns.

In particular, police leaders respond to these competing incentives by increasing the range of permissible tactics as officers become more biased toward aggressive tactics, giving those officers room to remain in compliance with policy even though their preferred (and chosen) tactics might not be

optimal from the perspective of producing high-quality outcomes. However, there is a limit to a police leader's willingness to engaging in this strategy. Once a police officer is too biased toward aggressive tactics, the police administrator's incentive to produce high-quality policing outcomes overrides the incentive to keep officers in compliance. Further, the implications of this logic shed light on three aspects of contemporary debate about the politics of public safety.

Establishing an effective police force

Policing is dangerous work that entails transferable skills. Officers may leave one jurisdiction for another or leave police work altogether, seeking employment in private security. Those facts create a challenge for leaders who need to recruit, train, and then retain competent and effective police officers. The model I develop illustrates one way in which structural forces can explain why jurisdictions throughout the United States are experiencing vacancy crises in their forces. Limited resources for police often mean low salaries and, even when police can be paid a competitive wage, political pressure and constraints on the job can easily offset the attractiveness of the profession. Consequentially, the extent to which the police officer's job becomes less attractive as wages decrease is related to the officer's underlying characteristics, especially the officer's bias with respect to how to handle himself in a policing event. Thus, there is a double-edged effect of political and resource constraints, in that it reduces the attractiveness of being a police officer for high-quality candidates, especially in places where the most serious policing events occur.

Democratic oversight and police accountability

The model I develop also provides insights into a complementary set of reforms that are at the heart of contemporary debate. Many American cities have either contemplated or adopted civilian oversight mechanisms, although the consensus is that they generally fail to bring about reform, owing to co-option by the police themselves and a lack of formal powers to induce changes in policy (e.g., Bobb 2003). One way in which that kind of oversight manifests is policy directives or blunt reforms adopted in response to perceived problems in policing practice, and that approach can be effective (e.g., Mummolo 2018b).

The model underscores an inherent trade-off between inducing compliance with policy while also giving officers the flexibility to respond to the inevitable extreme events that police must confront. That trade-off lies in the officer's temptation to violate policy. In this way, we see challenges a police chief can face in simultaneously satisfying the potentially competing pressures to keep officers in compliance with policy and produce high-quality policing. That tension

may help account for the pattern of short tenure for police chiefs across the United States.

Indeed, police leaders need to balance competing incentives for the police officers, and exogenously imposed constraints on the tactics and discretion of officers can have unintended consequences for the choices they make in particular settings. This insight does not mean that civilian or political oversight and accountability does not have its intended effect or produce better public safety. Rather, the model highlights an important interaction between policy and the conditions in which policing takes place, suggesting that involving police leaders in the reform process can enhance the extent to which reforms are well tailored to the particular needs of a police setting.

Police administration and public safety

Of course, police administration is just one aspect of public safety, and the model I have analyzed is limited in scope and only captures a small subset of the factors that influence both policy and police behavior. The model notably omits the public and its role in government; political leaders, in turn, are only mechanically modeled. Both the public and politicians are key actors in the establishment and provision of public safety, However, the model establishes a framework from which models of these additional features can build. Going forward, as scholars continue to examine the politics of police reform, the analysis presented here can provide guidance into the political economy of policy interventions designed to limit police excess.

In particular, contemporary challenges in the politics of public safety have brought about myriad calls for reform and questions about the relationship between state power and civil liberties. As political debates continue and politicians seek to identify better mechanisms for public safety in modern society, it will be important to consider carefully the nuances and complexities of police administration as they operate in the larger political context. Just as developing democracies often face challenges building state capacity and the rule of law, so too do established democracies have to consider the ways in which those same institutions evolve in social context.

ACKNOWLEDGMENTS

I thank Ethan Bueno de Mesquita, Adam Glynn, Andrew Little, Pablo Montagnes, Michael Leo Owens, John Patty, and Maggie Penn for helpful comments and discussions.

REFERENCES

- Alesina, Alberto, and Guido Tabellini. 2007. "Bureaucrats or Politicians?" Pt. 1, "A Single Policy Task." *American Economic Review* 97 (1): 169–79.
- Ashworth, Scott. 2012. "Electoral Accountability: Recent Theoretical and Empirical Work." *Annual Review of Political Science* 15:183–201.

- Balfe, Angela Rene. 2015. "The Revolving Door: A Correlational Study of Police Chiefs' Tenure." PhD diss., Marian University.
- Bass, Sandra. 2000. "Negotiating Change: Community Organizations and the Politics of Policing." *Urban Affairs Review* 36 (2): 148–77.
- Bendor, Jonathan, and Adam Meirowitz. 2004. "Spatial Models of Delegation." *American Political Science Review* 98 (2): 293–310.
- Bernstein, Maxine. 2019. "New Portland Police Recruits Failing Probation at Double the Rate of Past Years." *Oregonian*, March 13. <https://www.oregonlive.com/portland/2019/03/new-portland-police-recruits-failing-probation-at-double-the-rate-of-past-years.html>.
- Bies, Katherine J. 2017. "Let the Sunshine In: Illuminating the Powerful Role Police Unions Play in Shielding Officer Misconduct." *Stanford Law and Policy Review* 28:109.
- Bobb, Merrick. 2003. "Civilian Oversight of the Police in the United States." *Saint Louis University Public Law Review* 22 (1): 151–66.
- Carpenter, Daniel. 2020. *The Forging of Bureaucratic Autonomy: Reputations, Networks, and Policy Innovation in Executive Agencies, 1862–1928*, vol. 173. Princeton, NJ: Princeton University Press.
- Clark, Tom S., Elisha Cohen, Adam Glynn, Michael Leo Owens, Anna Gunderson, and Kaylyn Jackson Schiff. 2023. "Are Police Racially Biased in the Decision to Shoot?" *Journal of Politics* 85 (3): 826–42.
- Cohen, Elisha, Anna Gunderson, Kaylyn Jackson, Paul Zachary, Tom S. Clark, Adam N. Glynn, and Michael Leo Owens. 2019. "Do Officer-Involved Shootings Reduce Citizen Contact with Government?" *Journal of Politics* 81 (3): 1111–23.
- Davis, Robert C., Christopher W. Ortiz, Samantha Euler, and Lorriane Kuykendall. 2015. "Revisiting 'Measuring What Matters': Developing a Suite of Standardized Performance Measures for Policing." *Police Quarterly* 18 (4): 469–95.
- Dewan, Shaila. 2017. "Deconstructing the 'Ferguson Effect.'" *New York Times*, March 29. <https://www.nytimes.com/interactive/2017/us/politics/ferguson-effect.html>.
- Epstein, David, and Sharyn O'Halloran. 1994. "Administrative Procedures, Information, and Agency Discretion." *American Journal of Political Science* 38 (3): 697–722.
- Feuille, Peter, and Hervey A. Juris. 1976. "Police Professionalization and Police Unions." *Sociology of Work and Occupations* 3 (1): 88–113.
- Fisk, Catherine L., and L. Song Richardson. 2017. "Police Unions." *George Washington Law Review* 85 (3): 712–99.
- Gailmard, Sean, and John W. Patty. 2007. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51 (4): 873–89.
- Gailmard, Sean, and John W. Patty. 2012. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15:353–77.
- Geller, William A. 1985. *Police Leadership in America: Crisis and Opportunity*. New York: American Bar Foundation.
- Gelman, Andrew, Jeffrey Fagan, and Alex Kiss. 2007. "An Analysis of the New York City Police Department's 'Stop-and-Frisk' Policy in the Context of Claims of Racial Bias." *Journal of the American Statistical Association* 102 (479): 813–23.
- Gillion, Daniel Q. 2020. *The Loud Minority: Why Protests Matter in American Democracy*, vol. 20. Princeton, NJ: Princeton University Press.
- Gunderson, Anna, Elisha Cohen, Kaylyn Jackson Schiff, Tom S. Clark, Adam N. Glynn, and Michael Leo Owens. 2021. "Counterevidence of Crime-Reduction Effects from Federal Grants of Military Equipment to Local Police." *Nature Human Behaviour* 5:194–204.
- Hawdon, James E., John Ryan, and Sean P. Griffin. 2003. "Policing Tactics and Perceptions of Police Legitimacy." *Police Quarterly* 6 (4): 469–91.
- Hickman, Matthew J. 2006. "Citizen Complaints about Police Use of Force." US Department of Justice, Office of Justice Programs, Bureau of Justice.
- Hicks, Jerry. 1986. "Academy Tests Recruits' Resolve: 1 in 3 Fail to Finish 18 Weeks of Training." *Los Angeles Times*, December 12. <https://www.latimes.com/archives/la-xpm-1986-12-12-me-2399-story.html>.
- Huber, John D., and Nolan McCarty. 2004. "Bureaucratic Capacity, Delegation, and Political Reform." *American Political Science Review* 98 (3): 481–94.
- Huber, John D., and Charles R. Shipan. 2006. "Politics, Delegation, and Bureaucracy." In Robert Goodin, ed., *The Oxford Handbook of Political Science*. Oxford: Oxford University Press.
- Keenan, Kevin M., and Samuel Walker. 2004. "An Impediment to Police Accountability: An Analysis of Statutory Law Enforcement Officers' Bills of Rights." *Boston University Public Interest Law Journal* 14 (2): 185–250.
- Knox, Dean, Will Lowe, and Jonathan Mummolo. 2020. "Administrative Records Mask Racially Biased Policing." *American Political Science Review* 114 (3): 619–37.
- Lipsky, Michael. 1971. "Street-Level Bureaucracy and the Analysis of Urban Reform." *Urban Affairs Quarterly* 6 (4): 391–409.
- Lowande, Kenneth. 2020. "Police Demilitarization and Violent Crime." Working paper, University of Michigan.
- Maskin, Eric, and Jean Tirole. 2004. "The Politician and the Judge: Accountability in Government." *American Economic Review* 94 (4): 1034–54.
- Mastrofski, Stephen D. 2004. "Controlling Street-Level Police Discretion." *Annals of the American Academy of Political and Social Science* 593 (1): 100–118.
- Matusiak, Matthew C. 2016. "Dimensionality of Local Police Chiefs' Institutional Sovereigns." *Policing and Society* 26 (7): 753–70.
- McCall, Andrew. 2019. "Resident Assistance, Police Chief Learning, and the Persistence of Aggressive Policing Tactics in Black Neighborhoods." *Journal of Politics* 81 (3): 1133–42.
- McCubbins, Mathew D., Roger G. Noll, and Barry R. Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law, Economics, and Organization* 3 (2): 243–77.
- Meyer, Greg. 1992. "Nonlethal Weapons vs. Conventional Police Tactics: Assessing Injuries and Liabilities." *Police Chief* 59 (8): 10–18.
- Mummolo, Jonathan. 2018a. "Militarization Fails to Enhance Police Safety or Reduce Crime but May Harm Police Reputation." *Proceedings of the National Academy of Sciences* 115 (37): 9181–86.
- Mummolo, Jonathan. 2018b. "Modern Police Tactics, Police-Citizen Interactions, and the Prospects for Reform." *Journal of Politics* 80 (1): 1–15.
- Oates, Daniel. 2020. "I Used to Be a Police Chief: This Is Why It's So Hard to Fire Bad Cops." *Washington Post*, June 12.
- Rainguet, Fred W., and Mary Dodge. 2001. "The Problems of Police Chiefs: An Examination of the Issues in Tenure and Turnover." *Police Quarterly* 4 (3): 268–88.
- Schafer, Joseph A., Beth M. Huebner, and Timothy S. Bynum. 2003. "Citizen Perceptions of Police Services: Race, Neighborhood Context, and Community Policing." *Police Quarterly* 6 (4): 440–68.
- Sklansky, David A. 2008. *Democracy and the Police*. Redwood City, CA: Stanford University Press.
- Stephenson, Matthew C. 2007. "Bureaucratic Decision Costs and Endogenous Agency Expertise." *Journal of Law, Economics, and Organization* 23 (2): 469–98.
- Ting, Michael M. 2008. "Whistleblowing." *American Political Science Review* 102 (2): 249–67.
- US Department of Justice. 2018. "Full-Time Employees in Law Enforcement Agencies, 1997–2016." NCJ 251762, Office of Justice Programs, August.