Political Control and the Sources of Bureaucratic Change*

Tom S. Clark†

November 9, 2026

Abstract

Political control of bureaucracies operates through hierarchical mechanisms as well as environmental factors that shape the conditions under which decisions are made. While theoretical arguments illuminate these sources of change and the mechanisms through which they operate, the sources and mechanisms can be difficult or impossible to distinguish empirically. I propose a simple strategy for distinguishing electoral and environmental sources of change and the institutional mechanisms through which they operate. Using data from Cook County linking policing and prosecution records, I study the impact of the election of a reform prosecutor on how discretion is exercised in felony cases. I show that a policy adopted by the prosecutor produced a sharp, targeted response, while the change in discretionary decision-making began before the transition and occurred primarily among bureaucrats already in office. The results show how political control operates through incumbents and can begin before formal succession, with implications for studies of electoral turnover.

*I thank Jon Rogowski for helpful comments. I used Claude Opus 5.5 to assist in copyediting and formatting text as well as preparing some portions of the appendix. I used Codex GPT-5.6 Sol to prepare and debug code and create tables and figures.

†Professor of Political Science, Stanford University

1 Introduction

Prosecutors exercise tremendous power over American life, wielding nearly unfettered discretion over the application of criminal law. In the US, most prosecutors are elected (Gordon and Huber 2002), and those who are not nevertheless typically are directly accountable to elected politicians. Over the past decade, voters across the country have elected candidates promising to change how that power is exercised. In the federal system, concern has also grown about federal prosecutors' independence from political influence. Yet prosecutors influence the administration of judges from their position the the top of bureaucratic organizations in which individual cases are handled by line attorneys who exercise considerable (Lipsky 1980). This creates a familiar problem of political control: when does a change in political leadership alter the underlying bureaucracy, and through which institutional mechanism?

Past research has taught us a great deal about the tools elected officials use to hold bureaucracies in check and about when those tools bind (Gailmard and Patty 2012, McCubbins, Noll and Weingast 1987, Wood and Waterman 1991). However, we know less about what patterns of change in a bureaucracy reveal about those various mechanisms or how bureaucratic change can distinguish changes in leadership from changes in the broader political environment. In particular, a transition in leadership may be politically consequential without itself being the root cause or the temporal point of bureaucratic change.

Among political scientists, the dominant account identifies political leaders as the source of bureaucratic change. The principal-agent tradition highlights the oversight powers a leader has as a mechanism that can induce agents to change how they exercise discretion. A new leader can replace personnel or announce rules. She can also move the incumbents who remain through supervision and the prospect of reassignment (Bolton, de Figueiredo and Lewis 2021, Sasso 2020, Weingast and Moran 1983, Wood and Waterman 1991). An alternative account points to the environment in which bureaucrats exercise discretion. Scandals or incidents that lead to heightened scrutiny may alter what an office can afford to be seen

doing, even before a change in leadership (Busuioc and Lodge 2016, Carpenter 2010). From this perspective, leadership transitions may happen in the midst of bureaucratic changes that have already begun. Separating these processes can be difficult or even impossible because the same event may be causing both. Career officials may also respond to an anticipated successor. As a result, a change in who holds the leadership office may not trigger the onset of bureaucratic change.

Prosecutors are an unusually revealing setting in which to study this problem. They decide whom to charge and on what terms, their decisions are largely unreviewable, and in nearly every state they answer directly to voters rather than to an administrative superior (Gordon and Huber 2002). Recent studies use reform prosecutors taking office to estimate changes in charging and conviction (Agan et al. 2025) or imprisonment (Gottlieb et al. 2025). Others examine crime (Petersen, Mitchell and Yan 2024) or isolate the effects of marginal charging decisions (Agan, Doleac and Harvey 2023). These studies illuminate important consequences of prosecutorial reform. However, changes in criminal justice outcomes following electoral turnover do not necessarily isolate the effect of the elected prosecutor. Choices made by police officers, for example, determine which cases come to prosecutors in the first instance, so a charging rate reflects both what prosecutors decide and the cases police bring forward (Knox, Lowe and Mummolo 2020). The office may also have begun adapting before the prosecutor takes office.

I offer a simple framework that separates the source of changes in prosecution from the institutional mechanism through which that change occurs. In particular, changes in how prosecutors handle cases may stem from either the preferences of a current or anticipated prosecutor or changes in the office's reputational and political environment. The framework contemplates three mechanisms. The first is personnel, which involves a leader making choices about which line attorneys handle cases. Personnel change alters the distribution of stable reviewer types. In contrast, office-wide changes that all reviewers respond to will appear as a within-reviewer change that is common to all reviewers. A second is policy,

focusing on how new directives or rules can constrain discretion. Rules about which cases may or may not be prosecuted should produce a sharp response among covered cases. The third a change in the common standard reviewers apply, either due to direction from a leader or a response to the office's reputational environment. If any such change is simply due to a change in political leadership, then this kind of office-wide change in discretion should be associated with electoral events themselves (elections or inaugurations), whereas political anticipation implies an earlier beginning to the change. Even without isolating the political source of a change, distinguishing these mechanisms identifies how an office changed.

I evaluate these implications in the Cook County State's Attorney's Office from 2014 through 2019. Kim Foxx defeated the incumbent in the March 2016 Democratic primary and took office that December after campaigning to reform the office. I empirically examine the decision to approve or reject a felony arrest, which is known as felony review. These decisions are made by assistant state's attorneys roughly at the bottom of the professional hierarchy who are assigned to each arrest in a manner close to random conditional on the roster and call time (Jordan 2024). The analysis distinguishes variation in charging due to the distribution of reviewers from variation in how reviewers decide. Using a formal policy Foxx adopted under which the office would decline to prosecute retail theft under $1,000 as a felony, I benchmark discretionary changes in felony approval against formal constraints imposed by policy. I then link data on all arrests to felony charging decisions, to assess how much variation is driven by the cases police present to prosecutors as opposed to the decisions prosecutors make.

Over the period covered here, felony-review approval declined by about four percentage points, and most of that decline appears as common change within reviewers rather than a shift in the mix of reviewers. The office's policy against charging retail theft under $1,000 as a felony produced a sharp decrease in approval for that offense, but it does not explain the broader path. The broader decline was already under way before Foxx took office. There is no comparable break associated with either the Democratic primary or her inauguration,

although the estimates do not rule out gradual influence after office-taking. Changes in the case pool also do not account for the decline. The decline therefore appears primarily as a common trend among active reviewers rather than in who the attorneys are or in case composition. Of course, this is a study of just one office and so cannot identify the total causal effect of Foxx's election. More broadly, the results show that substantial bureaucratic change can precede formal succession and appear in the decisions of incumbent officials. Accounts of political control that begin with office-taking may therefore begin too late.

2 Political Change in a Bureaucracy

Like virtually every bureaucracy in government, political control of a prosecutor's office operates through delegation. In most American jurisdictions, voters select a prosecutor who in turn delegates to staff attorneys the responsibility for handling individual cases (Gordon and Huber 2002). In all but the smallest places, there are simply too many cases for the elected prosecutor to handle herself who are subject to a variety of standard agency mechanisms of oversight, including supervision and termination (e.g., Gailmard and Patty 2012, McCubbins, Noll and Weingast 1987, Weingast and Moran 1983, Wood and Waterman 1991). These tools can change who makes decisions, but they can also affect how attorneys use their discretion in the first instance. In the spirit of canonical principal-agent models, attorneys who expect closer supervision or reassignment have strong reasons to adapt, even though their own sincere preferences remain constant. In other words, political control can operate through incumbent officials rather than through the replacement of personnel.

The most immediate mechanism by which political principals affect bureaucratic decision-making is by inducing agents to use their discretion differently. And that need not align with the actual turnover in political control. Bureaucratic anticipation of a future principal can produce behavioral changes just as bureaucrats respond to incentives created by a current principal. This can take the form of a career official forming expectations about future

political principals during the course of an electoral cycle, and that anticipation can alter the official's decision-making even before the new principal possesses formal authority (Bolton, de Figueiredo and Lewis 2021, Sasso 2020). Of course, the incumbent may respond as well, facing an electoral threat that creates an incentive to change how the office operates before voters decide whether she will retain control. Further, even after an election that leads to a change in leadership, the new leader's impact may take a while to be felt as supervision and training alter how line officials understand the priorities of the office (Bowman, Lowrey-Kinberg and Gould 2026, Oberfield 2014). A new leader may therefore come into office in the midst of an institutional response already under way. Analyses focusing on the effects of turnover in an office may misstate the timing of political influence even when the incoming leader eventually matters.

But anticipation of turnover is only one reason an office may change before its leadership does. A bureaucracy may very well be sensitive to perceptions among the audiences that judge it. For example, scandals may increase the reputational cost of decisions perceived to be wrong or illegitimate, changing what officials regard as defensible without a directive from above (Birkland 1998, Carpenter 2010, Gilad, Maor and Ben-Nun Bloom 2015, Kingdon 1984). In the context of a prosecutor's office, that means line attorneys may exercise the same formal authority differently because the environment in which they exercise it has changed. When the leader is an elected official, the same event may also discredit an incumbent and alter the prospects of a challenger. Thus, political events may change both the office's reputational environment and its expectations about future leadership. Changes we observe in bureaucratic decision-making can therefore reflect either pressure—a distinction that timing alone does not make.

I distinguish these sources of change from the institutional mechanisms through which change occurs. Hierarchical pressure originates with the current or anticipated prosecutor. Environmental pressure originates with the audiences evaluating the office. The first mechanism is personnel composition: changing which attorneys exercise discretion (e.g., Lewis

2008, Wood and Waterman 1991). The second is policy, which typically takes the form of formal rules that constrain discretion for some set of cases (e.g., Huber and Shipan 2002, McCubbins, Noll and Weingast 1987). The third is office practice that affects the standard reviewers apply (e.g., Brehm and Gates 1997, Oberfield 2014). When office practice changes, incumbent reviewers use their discretion differently, without a new policy or a change in personnel. That change may reflect direction from a leader or a response to the office's political environment.

These mechanisms imply different forms of political control. If change occurs through personnel composition, then the prosecutor changes office behavior by changing which attorneys exercise discretion, rather than incumbent attorneys responding to supervision. If a prosecutor can bring about sharp changes via formal policy announcements, then perhaps she can control decisions in cases she identifies *ex ante*, but that does not inform the exercise of discretion outside the policy. By contrast, changes that operate through office practice change how prosecutors handle cases without being replaced or having each decision prescribed by policy. Political control then operates through discretion that prosecutors exercise, due to effective supervision or anticipated sanctions, in the spirit of standard principal-agent models. However, this kind of change may also reflect reputational pressure, so it does not by itself identify the source of change.

3 A Dynamic Model of Felony Review

I build a simple dynamic model that connects political pressure to charging decisions made by line lawyers working within an office led by a prosecutor. The model captures a prosecutorial office that comprises both an elected leader (the prosecutor) and career line attorneys (reviewers) who approve or reject felony charges subject to variation in both policy and practice. The modeling goal is to separate the effects of personnel, policy, and practice.

Players and characteristics. The model has an office and a set of reviewers, indexed by a and takes place in discrete time, $t = 1, 2, \ldots$, with cases indexed by i. I assume each case i is characterized by a degree of prosecutable strength $v_i \in \mathbb{R}$, an offense o_i, and characteristics x_i, which include the evidence, legal sufficiency, and other case-specific considerations relevant to felony charging. I assume v_i is drawn from a common knowledge distribution, conditional on x_i, with CDF$G(\cdot \mid x_i)$.

In each period t, there is a set $\mathcal{R}_t$ of active reviewers, whose membership may change from one period to the next, across whom the set of cases are distributed in each period. Each reviewer a has a preferred charging threshold $\alpha_a \in \mathbb{R}$, on the same scale as v_i, which is the minimum prosecutable strength she would require in the absence of any other considerations. The threshold is fixed for each reviewer but varies across reviewers. Let π_{at} denote reviewer a's share of cases in period t, with $\pi_{at} = 0$ when $a \notin \mathcal{R}_t$. Let $F_t(\alpha)$ denote the case-weighted distribution of preferred thresholds induced by these shares.

At each time t, there is an elected prosecutor in office whose intrinsic charging standard is $\bar{s}_t$. (I assume prosecutors differ only in this parameter.) There is also at each time a political state, Z_t, which contains the office's reputational and the electoral state. The electoral state includes the incumbent's position, the calendar, and beliefs about future control. I assume that in addition to her own intrinsic standard, the prosecutor also cares about the reputational component of Z_t, which I denote E_t, so that her ideal point at time t is

$$\hat{s}_t(E_t) = \bar{s}_t + \phi_E E_t, \qquad \phi_E > 0, \tag{1}$$

Choices and timing. At each time t, the office inherits the standard s^*_{t-1}, which is the standard implemented in the previous period. The incumbent observes the inherited standard and the political state, then chooses an office standard, $s_t \in \mathbb{R}$. Each reviewer $a \in \mathcal{R}_t$ observes that choice and chooses the charging threshold $\tau_{at} \in \mathbb{R}$ she will apply. Cases then arrive and are assigned to active reviewers. Let $D_{it} = 1$ if a formal policy in force at t

covers case i, and $D_{it} = 0$ otherwise. The assigned reviewer observes v_i and approves the case ($Y_{ait} = 1$) when its prosecutable strength exceeds her chosen threshold, unless a formal policy covers it:

$$Y_{ait} = \begin{cases} \mathbf{1}\{v_i > \tau_{at}\} & \text{if } D_{it} = 0, \\ 0 & \text{if } D_{it} = 1. \end{cases} \tag{2}$$

A formal policy leaves the reviewer's threshold unchanged but removes approval for the cases it covers.[1] The office standard becomes the inherited standard in period $t + 1$.

Preferences. Reviewer a prefers a threshold close to her preferred charging threshold but also bears a cost from departing from the office standard. She chooses τ_{at} to minimize

$$L_{at}(\tau_{at}; s_t) = (\tau_{at} - \alpha_a)^2 + \kappa\,(\tau_{at} - s_t)^2, \tag{3}$$

where $\kappa \geq 0$ is the weight she places on the office standard relative to her own judgment and captures the career and supervisory consequences of departing from the office standard.

The incumbent is forward-looking. In each period she incurs a loss

$$\ell_{\bar{s}_t}(s_t; s^*_{t-1}, E_t) = \big(s_t - \hat{s}_t(E_t)\big)^2 + \zeta(s_t - s^*_{t-1})^2, \qquad \zeta \geq 0, \tag{4}$$

where the second term captures the notion that changing the standard in any period is costly. I assume the incumbent discounts the future by $\delta \in (0, 1)$. Letting $V_{\bar{s}}(s, Z)$ denote the value to a prosecutor with intrinsic standard $\bar{s}$ when the inherited standard is s and the political state is Z, the incumbent's Bellman equation is

$$V_{\bar{s}}(s, Z) = \max_{s'} \left\{ -\ell_{\bar{s}}(s'; s, E) + \delta\, \mathbb{E}\left[\mathbf{1}\{\text{retained}\} V_{\bar{s}}(s', Z') + \mathbf{1}\{\text{replaced}\} \overline{V}_{\bar{s}} \mid s, Z, s' \right] \right\}, \tag{5}$$

[1] Of course, a formal policy requiring charging a case is theoretically feasible under some circumstances, but it is outside of the model I study here and would entail a number of complications and constraints related to ethical and legal guidelines and rules.

where $\overline{V}_{\bar{s}}$ is the incumbent's value after leaving office. The political state evolves according to a Markov transition that may depend on the standard the office implements. Outside an election period, the incumbent is retained with probability one. Whether the prosecutor is replaced or retained is determined by the transition. An incumbent can therefore change the standard before the election because doing so affects her prospect of remaining in office. A successor inherits the standard even though the departing incumbent does not value the successor's policy payoff.

Equilibrium. To solve the model, I characterize a Markov-perfect Stackelberg equilibrium. That means the incumbent's strategy depends on the inherited standard and the current political state, while reviewers observe the office's standard and best respond. (Full model details, formal reuslts, and proofs are reserved to Appendix A.)

In equilibrium, each reviewer applies a weighted average of her own preferred threshold and the office standard,

$$\tau_{at}^{*} = (1-w)\alpha_a + w s_t^{*}, \qquad w \equiv \frac{\kappa}{1+\kappa}, \tag{6}$$

where $s_t^{*} = \sigma_{\bar{s}_t}(s_{t-1}^{*}, Z_t)$ is the standard that solves the incumbent's decision problem given by Equation (5) and w is a weight that increases in the cost of departing from the office standard. The office standard, s_t^{*}, balances the incumbent's current ideal point against the cost of changing the inherited standard, the value of retaining office, and expectations about the political state that will follow, via her continuation value.

Composition and the office standard. The model distinguishes a change in who reviews cases from a change in the standard they apply. For a case not covered by a formal policy ($D_{it} = 0$), the probability the reviewers approves is given by

$$\mathcal{A}(\alpha, s_t^{*} \mid x_i) = 1 - G\big((1-w)\alpha + w s_t^{*} \mid x_i\big). \tag{7}$$

Thus, for a given case, approval is less likely when a stricter reviewer handles the case or when the office raises its standard. The threshold itself is linear in α and s_t^*, but the approval probability need not be, because I have not specified the distribution of prosecutable strength. To derives an empirical prediction, I assume the probability is linear around the charging margin. In particular, let μ denote approval at a reference threshold and case mix, and let $\bar{f} > 0$ denote the density of prosecutable strength at that margin.[2] To first order approximation,

$$\mathcal{A}(\alpha, s_t^* \mid x_i) \approx \mu - \bar{f}(1-w)\alpha - \bar{f}ws_t^*. \tag{8}$$

Given this approximation both the reviewer's own threshold ($^*\alpha$) and the office standard (s_t^*) affect the probability linearly. Let $\widetilde{\alpha}_a = -\bar{f}(1-w)\alpha_a$ denote the effect of reviewer a's threshold, and $g(t) = -\bar{f}ws_t^*$ denote the the office standard (i.e., policy) component. Because α_a is fixed, $\widetilde{\alpha}_a$ is a stable reviewer type. Because s_t^* is the same for all reviewers, $g(t)$ shifts every reviewer's approval probability by the same amount. Averaging over reviewers with case shares π_{at} gives office-wide approval, $A_t \approx \mu + \sum_a \pi_{at}\widetilde{\alpha}_a + g(t)$, and the change in the review probability time time 0 to time t has two separable components:

$$\Delta A_t = \underbrace{\sum_a (\pi_{at} - \pi_{a0})\widetilde{\alpha}_a}_{C_t^F} + \underbrace{g(t) - g(0)}_{C_t^s}. \tag{9}$$

The first term changes as cases are assigned to reviewers with different thresholds, while the second term changes as the office standard changes, and this forms the basis of the empirical analysis below.

Anticipated supervision. Finally, consider a simple extension to the model's logic. In the model, reviewers respond to the incumbent's standatd, s_t^*. The incumbent may adjust that

[2]The approximation treats the density at the margin as common across cases, which conditioning on case characteristics accomplishes in the empirical model. It holds exactly when $G(\cdot \mid x_i)$ is linear in a neighborhood of the margin.

standard in anticipation of an election, but there is no incentive for reviewers to anticipate policy after the election. As I have emphasized, though, career bureaucrats often do look to the future in anticipation of a new political principal (Bolton, de Figueiredo and Lewis 2021, Sasso 2020). To allow for that dynamic in a simple way, let B_t denote the effects (through supervision or pressure, etc.) a reviewer anticipates from a future principal, which can depend on beliefs about future control in Z_t, and suppose a response to it shifts the standard reviewers apply by $\eta_B B_t$, with $\eta_B \geq 0$. That leads to an *extended* office standard given by

$$\widetilde{s}_t^* = s_t^* + \eta_B B_t, \tag{10}$$

and the reviewer's threshold becomes $\widetilde{\tau}_{at}^* = (1-w)\alpha_a + w\widetilde{s}_t^*$. This extension does not affect the decomposition in Equation (9), but it does add a second path by which the common component can move before a transition in leadership. However, because $g(t)$ depends on s_t^* and $\eta_B B_t$ only through their sum, the extension does not provide any further identification of the mechanisms.[3]

Predictions. These results imply that the probability a reviewer approves felony charges in a case not otherwise covered by a formal declination policy is simply $\Pr(Y_{ait} = 1 \mid x_i, a, t, D_{it} = 0) = 1 - G(\widetilde{\tau}_{at}^* \mid x_i)$. The first thing to notice is that the reviewer's identity and the office standard yield different empirical implications, as we saw in the decomposition above (Equation 9). Thus, a key prediction is that changes in felony review approval due to personnel mean that reviewer identity will explain variation in approval, whereas changes in practice (i.e., the office standard) will appear as a temporal pattern across all reviewers

Implication 1 (Composition and common movement) *Personnel change makes C_t^F dominate. A changing common standard makes C_t^s dominate and lowers approval among reviewers observed across the transition.*

[3]Common forces outside Z_t that moved over the period enter $g(t)$ as well, which weakens any attribution of its movement to the sources in the model.

Second, a formal policy the office standard have different implications for approval. A policy that removes discretion over some set of cases, whereas a change in the office standard shifts the threshold for cases left to reviewer discretion. To benchmark the effect of policy, I rely on the retail-theft declination policy. Because that policy should result in a sharp change in approval that is limited to the cases it covers, the imposition of the policy allows me to establish both the magnitude and the scope of such a policy-driven change.

Implication 2 (Categorical rule) *A formal declination rule produces a sharp change among cases that satisfy its eligibility conditions. If eligibility is concentrated in an observed offense and the composition of cases reaching review remains stable, that offense shows a differential decline at the rule's effective date. A changing office standard affects discretionary decisions beyond that offense.*

Finally, notice that an abrupt change in approval does not by itself identify the effect of a formal policy. An event that changes control of the office or beliefs about future control beyond prior expectations can also shift the office standard immediately. But, because I assume that changing office standards is costly, a change in the incumbent's identity is also consistent with a gradual change in her standard. On the other hand, an anticipated transition can affect $\widetilde{s}_t^*$ earlier through the incumbent's policy or reviewers' responses to expected supervision. The incumbent may also change the standard before the election when office conduct affects her probability of remaining in office.

Implication 3 (Onset) *An electoral event that changes control or beliefs beyond prior expectations produces a level shift or a change in the subsequent path. The absence of such an effect does not reject electoral influence operating through the incumbent's earlier adaptation or anticipation of a successor or environmental effects already under way.*

These implications organize the empirical analysis that follows. The analysis below estimates whether a decline in felony review approval rates is due to changes in reviewer composition or instead an office-wide change common to all felony-review attorneys. I also assess whether the response to the retail-theft rule was concentrated among the narrow set of cases it technically covered or instead changed broader practices in felony review. Finally, I evaluate whether political events related to the state's attorney's office changed the level

or subsequent trajectory of approval. Together, these tests can identify the organizational mechanism and timing of change, but they cannot (by themselves) determine what caused any change in the office standard. The extended office standard, $\widetilde{s}_t^* = \sigma_{\bar{s}_t}(s_{t-1}^*, Z_t) + \eta_B B_t$, combines the incumbent's dynamic response to reputational and electoral conditions with reviewers' common response to a current or anticipated principal. For example, the McDonald video changed the reputational environment and the incumbent's electoral position at the same time. Foxx's candidacy also changed beliefs about future control. These events, though, might collectively have felony approval in a similar way, and so $\widetilde{s}_t^*$ cannot separate their effects.[4]

4 Data and Empirical Strategy

I evaluate the model's predictions in the context of felony review in the Cook County State's Attorney's Office (CCSAO). In Cook County a typical felony case involves the police making an arrest that is then subject to approval by the CCSAO.[5] Specifically, after making a felony arrest, the police call the CCSAO's Felony Review Unit, where an on-call assistant state's attorney usually evaluates the arrest within twenty-four hours. The reviewer approves felony charges, rejects them, or requests further investigation. Reviewers work fixed shifts, and incoming calls are assigned in rotation to the next reviewer on the active roster. Conditional on the roster and call time, the reviewer assigned to a case is therefore close to random with respect to its characteristics (Jordan 2024). The reviewer handles only this stage and then

[4]Appendix A describes the source-specific heterogeneity that could distinguish these sources. The available data do not contain the predetermined reviewer or offense characteristics needed for that test, so I treat the source of the common movement as unresolved.

[5]Illinois law permits police to file charges directly in narcotics cases, assuming no other felony is involved (Jordan 2025), and a recent experimental policy has expanded that authority to gun charges. Narcotics-only arrests account for roughly half of Chicago's felony arrests during the period, so I exclude them from the population at risk of reaching review.

turns the case over to another attorney. Most reviewers are junior attorneys who rotate through the unit within a year or two.

The data here cover the period from January 2014 through December 2019. I chose this period because it spans a police scandal and a transition in leadership. In particular, in November 2015 the Chicago Police Department released a video of the police killing of Laquan McDonald a year earlier, following a court order to do so. The video contradicted official accounts of the killing and was central to the subsequent murder prosecution of the officer who shot him. Shortly thereafter, in March 2016, Kim Foxx defeated the incumbent in the Democratic primary. She later won the general election and took office on December 1, 2016. Of course, the scandal had direct electoral implications, because the video discredited the incumbent and improved Foxx's prospects. Two weeks after coming to office, Foxx's administration adopted a policy under which the office would not charge individuals with felony retail theft unless the amount stolen exceeded $1,000, despite state law defining felony retail theft as stealing merchandise valued above $300.

Specifically, I build a case-participant panel that follows individual cases from arrest through charging. I identify the universe of arrests using data from the Chicago Police Department, which includes the charges, governing statutes and offense classes, arrestee race, and arrest date. I link these records to the State's Attorney's databases that cover felony intake and each charge initiated against each arrestee. Using data from an Illinois Freedom of Information Act request, I link these data sources via the booking number and the State's Attorney's case-participant identifier. A separate crosswalk connects that identifier to the public case records. These links allow me to follow a Chicago arrest through felony review and charging. Roughly 95 percent of felony arrests match to a case record, with no meaningful trend from 2014 through 2018.[6]

Because of data availability, the arrest data include only arrests made by the Chicago Po-

[6]The match rate declines modestly in late 2018, which I address in Appendix G. Unmatched records include both genuine non-prosecution and errors in booking numbers, which I cannot fully separate.

Table 1: Sample description, by period. Countywide reviewed case-participants, 2014–2019.

	Pre (2014–16)	Post (2017–19)
Reviewed case-participants	59,813	51,031
Black (%)	61.0	66.3
White (%)	15.8	12.2
Hispanic (%)	20.9	19.4
Other (%)	2.3	2.2
Male (%)	85.4	87.4
Mean age at incident	33.2	32.8
Felony-class charge (%)	99.8	99.9
Felony-review approval (%)	97.8	94.0

Columns pool the 2014–16 and 2017–19 periods. These records cover all cases from Cook County. The data linked to arrests in Figure 1 include only Chicago Police Department arrests.

lice Department, the county's dominant source of felony cases. Other municipal departments and the county sheriff also send cases to the State's Attorney. Arrests from 249 agencies enter felony review during the study period, and 67.4% come from the Chicago Police Department. As a result, for the analyses that use only review and charging records, I can use the full countywide caseload. Because the arrest data include arrestees' race but not age or sex, arrest-side composition adjustments use race, offense, and charge class. Table 1 describes the countywide reviewed caseload before and after the transition. Figure 1 shows the full data funnel for all Chicago arrests I was able to link to prosecution data.

4.1 Measures

With these data, I build a set of variables on which the empirical strategy relies.

Felony-review approval. This variable captures the central charging-stage outcome for each case-participant, measured at the level of the primary charge. I code a case as approved if its felony-review result is recorded as "Approved" or "Charge(S) Approved," and as not approved if it is "Disregard," "Rejected," or "Continued Investigation." I exclude observations where the felony review outcome is either missing or takes on an "other" value.

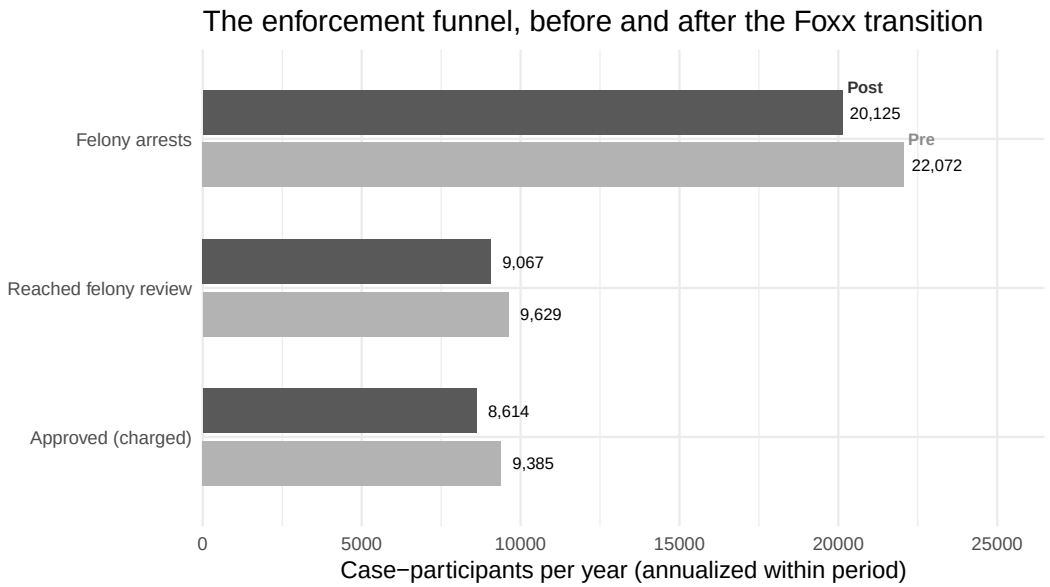


Figure 1: The enforcement funnel before (Pre, 2014–16) and after (Post, 2017–19) the Foxx transition: felony arrests, arrests reaching felony review, and approvals (charges filed). I calculate counts annually using the linked records from the Cook County State's Attorney and the Chicago Police arrest data, so each stage is a subset of the one above it.

Felony-eligible arrest and reaching review. For the analysis that involves selection into felony review, I need the population of felony-eligible arrests before the felony review stage. I code an arrest as felony-eligible if any of its charges is recorded as a felony, and as *review-eligible* if at least one of its felony charges is not a narcotics offense, since narcotics-only arrests bypass review by rule.[7] I code an arrest as having reached felony review if it appears in the felony review database. The selection rate is the share of review-eligible arrests that reach review, conditioned on a successful arrest-to-case link.[8]

Filed charge class. I order Illinois offense classes from most to least severe: felony classes X, 1, 2, 3, and 4, with first-degree murder (Class M) ranked alongside X, followed by misdemeanor classes A, B, and C. I use the class of the charge as filed at initiation to measure

[7]Narcotics statutes are 720 ILCS 570 (controlled substances), 550 (cannabis), 646 (methamphetamine), and 600 (drug paraphernalia).

[8]Conditioning on a successful link removes both cases mechanically excluded and the modest decline in match quality from the estimates.

charge severity. Illinois codes all felonies into one of these classes and maps each offense class into a sentencing range that constrains judges' decisions.

Reviewer identity. To distinguish changes in *who* reviews cases from changes in *how* reviewers decide, I estimate reviewer-specific fixed effects. These account for stable approval propensities, net of the common trends over time, the retail-theft policy, and observable case mix. The role of variation in attorneys (i.e., personnel) is the change in the case-weighted mean of these effects. Appendix D reports simpler leave-one-out and turnover measures as alternative diagnostics.

4.2 Empirical Strategy

The main empirical problem is that the cases prosecuted come from a population of arrests that the police, rather than the prosecutor, choose. Responding at least partially to their own broader political and institutional environment, police make the critical decision about which arrests to make and present to prosecutors in the first instance (Richman 2003, Shjarback et al. 2017, Sklansky 2018). Thus, changes in aggregate charging rates may reflect changes in the standards the office uses to evaluate cases or changes in the mix of cases presented to the office (Knox, Lowe and Mummolo 2020, Ouss and Rappaport 2020). Using the universe of arrests lets me observe the population before felony review, so I can describe how the case pool changed and bound whether that change can account for the movement in approval.[9] I present the analysis in three steps, corresponding to Implications 1, 2, and 3. The first step also asks whether changes in the reviewed caseload can account for the estimated common movement.

[9] I do not identify why police behavior changed or whether it responded to the prosecutor, a subject that deserves more theoretical and empirical attention (cf. Ouss and Rappaport 2020, Richman 2003).

The empirical analysis rests on a simple baseline specification:

$$\mathbb{E}[\text{approve}_{it} \mid x_{it}, a(i), t] = g(t) + \gamma\,\text{policy}_{it} + x_{it}'\beta + \underbrace{\widetilde{\alpha}_{a(i)}}_{\text{type}} + \underbrace{\nu_{a(i)}\,t}_{\text{within-reviewer deviation}} . \tag{11}$$

Here, I model a smooth time trend, $g(\cdot)$, to capture common changes over time across all reviewers. The variable policy_{it} indicates whether case i is a retail theft case covered by the retail theft policy at time t, and x_{it} contains case covariates. I estimate a fixed effect for each reviewer a with $\widetilde{\alpha}_a$. Given Equation (8), we can interpret this specification as a model of the reviewers' thresholds for felony review, which includes a variable reviewer-wide standard, $g(t) = -\bar{f}w\widetilde{s}_t^*$. The reviewer effect is the approval-scale reviewer type $\widetilde{\alpha}_a = -\bar{f}(1-w)\alpha_a$, which the model holds fixed over time. The policy term stands in for the categorical rule D_{it}. The reviewer-specific slope ν_a allows a reviewer to depart from the office-wide change in review thresholds.[10]

I estimate a series of variants of Equation (11). First, I use Equation (9) to separate changes in review rates into reviewer composition and change in the common standard. I then examine variation in approval among reviewers observed before and after the leadership transition. Second, I evaluate the December 15, 2016 retail-theft policy to estimate how a binding policy changes prosecutorial discretion. This allows me to distinguish the effect of a formal rule from an office-wide change that occurs when Foxx comes to office. I then turn to Implication 3, which concerns the shape of $g(t)$ at key political moments. I estimate the model to assess whether there are shifts in either the rate of approval or its trend associated with the critical events discussed above. Of course, even if there is no break, it does not mean that mean electoral influence does not shape decision-making. It is possible that electoral influence operates by an incumbent shifting the office standard (i.e., practice) in advance of

[10]Appendix B reports the institutional basis for the assignment claim and balance tests. Assignment is determined the list of felony-review assistant state's attorneys on duty at the time an arrest is reported to the CCSAO, and observed case attributes explain little of which reviewer receives a case.

an election, through felony-review attorneys anticipation of turnover in the prosecutor, or through via incremental changes too small to empirically distinguish from a change already under way. Thus, I focus on what mechanisms we can rule out.

5 Estimating the Nature of Felony Review Changes

5.1 Implication 1: Personnel versus the common standard

The first empirical question I address is whether the office changed because different reviewers handled its cases or because active reviewers changed how they decided them. The same aggregate change in felony review might occur under under either mechanism, even for different reasons. Implication 1 identifies how to distinguish them in the data. Personnel change shifts cases toward reviewers with different stable approval propensities, making C_t^F the dominant term in Equation (9). The common standard mechanism changes approval within the active set of reviewers, making C_t^s the dominant term.

For the decomposition, I set $\nu_a = 0$ and estimate the following specification:

$$\mathbb{E}[\text{approve}_{it} \mid a(i), t, x_{it}] = g(t) + \gamma\, \text{policy}_{it} + x_{it}'\beta + \widetilde{\alpha}_{a(i)}, \tag{12}$$

where $g(t) = \beta_0 + \sum_{k=1}^{12} \theta_k B_k(t)$ and $B_k(t)$ are the twelve basis functions of a natural spline in review month. The policy indicator is

$$\text{policy}_{it} = \mathbf{1}\{o_i = \text{Retail Theft}\}\mathbf{1}\{t \geq \text{January 2017}\}.$$

The offense indicator captures whether the retail-theft rule applies to a given case (D_{it}). The vector x_{it} contains linear controls for a defendant's age and and the severity of the charge, as well as fixed effects for race, sex, and broad offense type. The parameter $\widetilde{\alpha}_{a(i)}$ captures the effect of the reviewing attorney. This specification is Equation (11) with $\nu_a = 0$ and $g(t)$ represented by the natural spline.

The reviewer fixed effects capture stable differences in approval propensity across reviewers. The personnel contribution, C_t^F, is the change between 2014 and 2019 in the case-weighted mean of those effects. This allows for a change in the share of cases handled by each reviewer while holding each reviewer's baseline approval propensity fixed. The common contribution, C_t^s, is the change in $g(t)$ over the same period and represents office-wide trends over time in the felony approval rate. I estimate the model using 93,378 decisions by 321 reviewers who handled at least 25 cases.

Recall from Table 1 that felony-review approval fell by about four percentage points over the window, from an average of 97.8 percent in 2014–16 to 94.0 percent in 2017–19. Figure 2 shows the estimate from Equation (12), which isolates the monthly trend net of the effect of the retail theft policy, reviewer-level effects, and case characteristics. Each gray point is the monthly approval rate after removing the estimated contributions of reviewer identity, the retail-theft policy, and observed case mix. The solid line is the fitted spline under the same normalization.[11]

The fitted line falls 4.27 percentage points between January 2014 and December 2019, dropping from 98.21% to 93.95%. Table 2 compares case-weighted annual averages; by that measure, the common component is 3.29 percentage points lower in 2019 than in 2014. Of course, the temporal component, $g(t)$, only captures a common trend across all reviewers and so ay be due to either policies implemented by an the incumbent state's attorney at the time or shared expectations about future standards, so these estimates cannot identify which mechanism is driving the change.

To do so, I break down the various estimates of the sources of change in Table 2. Because cases are essentially randomly assigned to reviewers conditional on date and time,

[11]Specifically, each point is the fitted spline for that month plus the mean regression residual among cases reviewed during the month. I shift both series by a single constant so that their case-weighted means equal the sample approval rate. The absolute vertical level is therefore a normalization; only differences in $g(t)$ over time are identified.

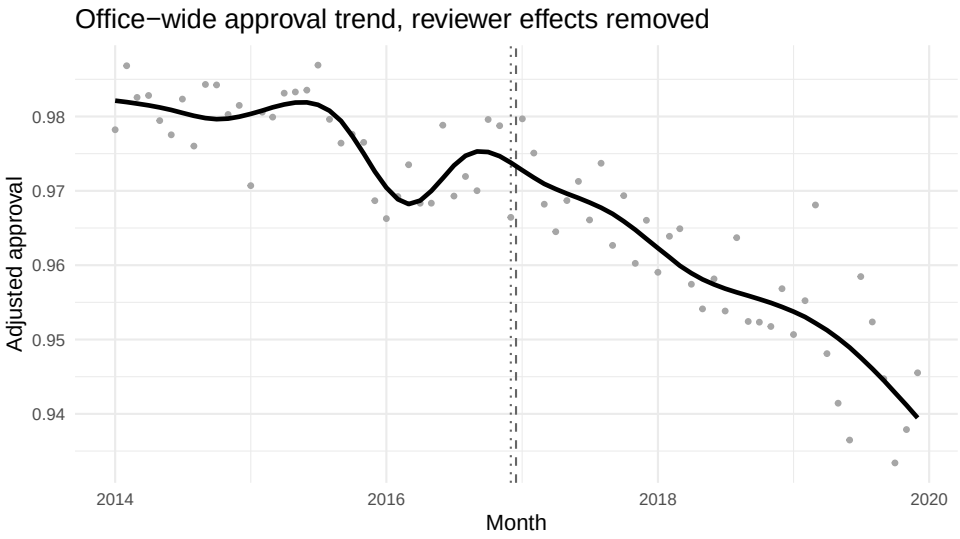


Figure 2: The office-wide calendar component of felony-review approval. The gray point for month t is $\widehat{\kappa} + \widehat{g}(t) + \overline{\widehat{\epsilon}}_t$, where $\overline{\widehat{\epsilon}}_t$ is the mean residual from the case-level model described in the text. The solid line is $\widehat{\kappa} + \widehat{g}(t)$. The constant $\widehat{\kappa}$ returns the series to the approval-rate scale; its level is a normalization. The dotted and dashed lines show Foxx's inauguration the adoption of the retail-theft rule.

the estimated reviewer effects can be interpreted as differences in reviewer type rather than differences in the cases assigned to them (see also Jordan 2024). Holding case mix and the retail-theft policy constant, approval declined by roughly 4 percentage points between 2014 and 2019. Reviewer composition accounts for very little of that decline (less than 1 percentage point), though the estimate is fairly imprecise. The office-wide time trend accounts for more than 3 percentage points, which is more than 80% of the decline. The point estimate therefore assigns about four times as much of the decline to common movement as to reviewer composition.

Turning from the office-wide, common change in approval, I evaluate change in approval among a stable group of reviewers. Ideally, I would estimate the distribution of the reviewer-specific slopes ν_a. Unfortunately, reviewer tenures are too short and uneven to estimate that distribution precisely. I instead restrict the sample to reviewers who made at least ten decisions on each side of the December 2016 inauguration and estimate annual changes relative to 2016. This holds the membership of the group fixed by construction. Among

Table 2: Additive decomposition of the change in felony-review approval.

Quantity	Estimate	95% CI
Additive OLS decomposition, 2014–2019		
Adjusted approval change (percentage points)	−4.05	[−4.58, −3.51]
Reviewer composition C^F (percentage points)	−0.76	[−2.08, 0.57]
Common movement C^s (percentage points)	−3.29	[−4.69, −1.89]
Composition share (percent)	18.7	[−13.9, 51.3]
Common-movement share (percent)	81.3	[48.7, 113.9]

The table implements Equation (9) in a linear probability model using reviewers with at least 25 cases. Reviewer types are fixed effects, and a twelve-degree-of-freedom natural spline measures the common time component. The 2014 and 2019 components use the actual case weights in each year, with inactive reviewers assigned zero weight. Confidence intervals cluster by review month. The reviewer–month graph forms a single connected component. Appendix D reports sensitivity to the caseload threshold and flexibility of the common trend.

the 82 reviewers who meet that requirement, adjusted approval in 2019 is 3.61 percentage points ([−6.87, −0.35]) below its 2016 level (see Figure 3). Requiring between five and twenty decisions on each side produces declines between 3.45 and 3.90 points. However, it is important to underscore that this does not imply that every reviewer moved in parallel. Rather, it simply shows that the decline was not confined to replacing reviewers with less permissive ones.

Taken together, the point estimates attribute most of the decline to a change in the office standard rather than reviewer composition. Critically, the decline in the rate of approval was not due simply to turnover in who serves as a felony-review attorney. Of course, the estimate cannot identify the source of the change and are consistent with the incumbent changing the office standard in advance of the election, reviewers anticipating her successor, or a change in the office's reputational environment.

Of course, the cases reaching felony review changed over the period. Chicago felony arrests fell from about twenty-five thousand in 2014 to about eighteen thousand in 2016–17, then rebounded to about twenty-two thousand by 2019. Among review-eligible arrests, the share reaching review fell from about 0.97 in 2014 to about 0.90 in 2019. Recalculating each year's approval rate using the 2014 proportions of cases in each race-by-felony-class category produces nearly the same annual trend. Moreover, while the approval rate fell over the

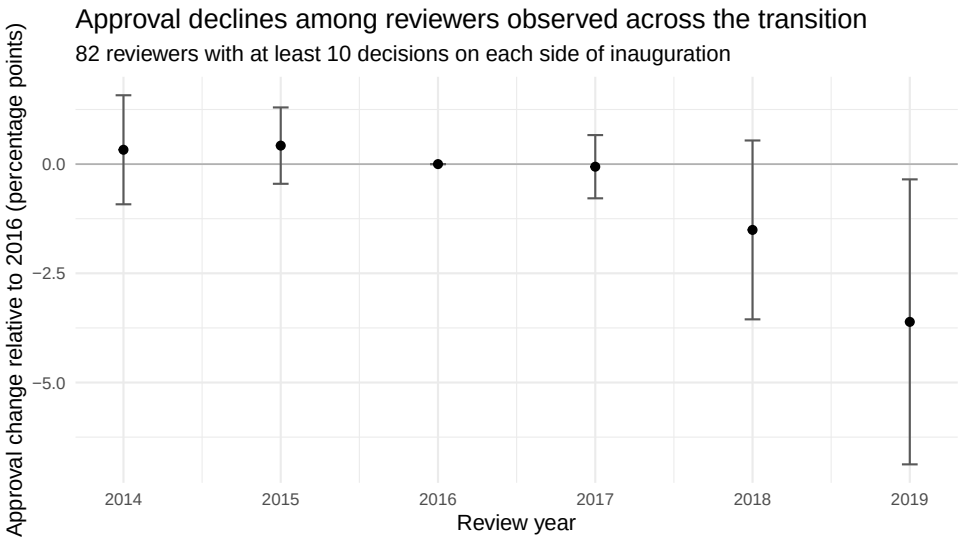


Figure 3: Annual approval changes among reviewers with at least ten decisions on each side of the December 2016 inauguration. Estimates are relative to 2016 and come from a linear probability model with reviewer fixed effects, case controls, and the retail-theft policy indicator. Bars report 95 percent confidence intervals clustered by reviewer.

period, tighter selection into review should, if anything, leave the office with stronger cases and push approval upward. Under the monotonicity assumption described in Appendix G, trimming bounds place the within-pool change below zero from 2016 forward. The covariate-tightened bound for 2018 is approximately $[-0.036, -0.027]$, and its Imbens–Manski (Imbens and Manski 2004) confidence interval also remains below zero. Collectively, these robustness checks suggest that the decline in approval is not due to changes in the cases presented to the office.

5.2 Implication 2: Categorical policy

Turning to Implication 2, the question is whether I can distinguish the effects of a categorical rule affecting felony review for some cases from an office-wide change in review thresholds. On December 1, 2016, Foxx took office and, two weeks later, announced a formal policy governing felony retail-theft declinations. Illinois law at the time classified retail theft above $300 as a felony, while ordinary theft carried a $500 threshold. Effective December 15, 2016, the office presumptively declined to charge retail theft as a felony unless the value reached

$1,000 or the accused had at least ten prior felony convictions. However, because the policy was announced just two weeks after Foxx took office, any change in felony-review approval rates around that policy could combine a targeted policy response with office-wide movement. Thus, I rely on the logic from Implication 2, which highlights that the rule should produce an immediate response, but only among the cases it covers.

Unfortunately, the data do not consistently report the value of the stolen property or the defendant's criminal history, so I cannot identify every case for which the rule bound (i.e., retail theft between $300 and $1,000). I instead estimate an offense-level intention-to-treat response by comparing all retail-theft cases with ordinary-theft cases. This comparison mixes retail-theft cases covered by the policy with cases that remained eligible for felony prosecution. Thus, the estimates I report here reflect both the share exposed and compliance with the rule. However, because the policy also changed which cases reached review, it need not be an attenuated version of the effect among covered cases. (It is possible, for example, that the policy simply stopped referring cases for review when they were covered by the declination policy.) In this sense, the estimates here constitute a reduced-form response for retail theft as an offense category rather than the structural effect of $D_{it} = 1$.

I estimate a simple event-study specification that modifies the baseline model. This specification allows the retail-theft differential to vary across periods. It shows whether the offenses were already diverging before the rule and whether the differential appears when the rule takes effect. Formally, I replace $g(t)$ with calendar-month fixed effects, restrict the sample to retail and ordinary theft, and estimate

$$\mathbb{E}[\text{approve}_{it} \mid a(i), t, x_{it}] = \mu_t + \rho R_i + \sum_{q \neq -1} \gamma_q R_i \mathbf{1}\{Q_{it} = q\} + x_{it}'\beta + \widetilde{\alpha}_{a(i)}, \tag{13}$$

where μ_t is a calendar-month fixed effect, R_i indicates retail theft, and x_{it} contains linear controls for age and charge severity and fixed effects for race and sex. The reviewer effect is $\widetilde{\alpha}_{a(i)}$. September through November 2016 form the omitted period, $q = -1$. December

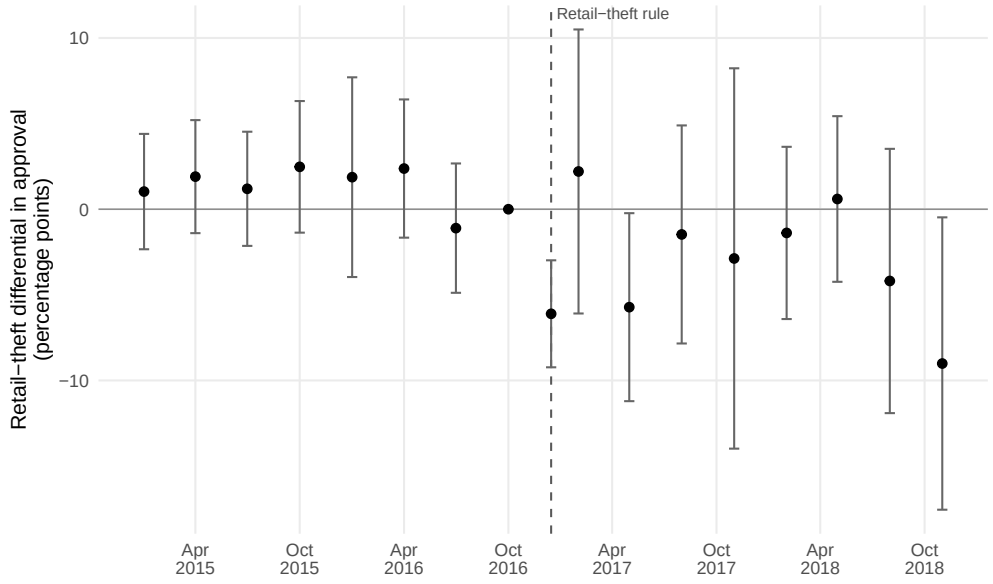


Figure 4: Retail-theft approval around the December 15, 2016 declination policy, relative to ordinary theft. Points are differences from the September–November 2016 retail-theft differential. December 2016 is its own transition period; the remaining estimates pool three months. Bars report 95 percent confidence intervals clustered by calendar month.

2016 is a separate transition period, $q = 0$, because the rule took effect halfway through the month. The remaining periods pool three months each, and the estimation window covers two years on either side of the policy. Standard errors are clustered by calendar month.

Figure 4 reports the estimates covering two years on either side of the policy. It groups months into three-month periods, except that December 2016 remains a separate transition period because the rule took effect halfway through the month. September through November 2016 form the omitted reference period. Before the adoption of the declination policy, retail and ordinary theft follow similar paths: the seven prepolicy coefficients are jointly indistinguishable from zero ($p = .56$). In December, retail-theft approval falls by 6.11 percentage points relative to ordinary theft and the prepolicy difference. Thus, the response to the policy takes place immediately after it is announced and is limited the the specific offenses it covers, rather than spilling over to other areas of felony review.

Pooling the full post-policy period produces the same conclusion: retail-theft approval fell by 3.56 ($[-5.77, -1.34]$). Appendix E reports the full pooled specification using the

December 15 treatment date, an alternative beginning in January 2017, and estimates using broader comparison groups. The December drop does not persist in every subsequent quarter, but the set of retail-theft cases reaching review also changes sharply. Specifically, the share of the felony review docket made up of retail theft cases fell from about 15% to about 5% within weeks and remained at that level through 2019 (Appendix G). The cases observed after the policy was adopted therefore need not be comparable to the pre-policy cases.

The key finding here is the contrast between the change in approval of cases (potentially) covered by the declination policy and the temporal trend in approval. Whereas the policy leads to an immediate change for the offense it names, the office-wide path remains smooth after holding that response apart. In this way, the results here distinguish an immediate, offense-specific change (of this magnitude) from an office-wide change in the common standard.

5.3 Implication 3: Change before political turnover

Finally, I turn to Implication 3, which concerns when the office-wide change in approval began. As the model illustrates, a variety of potential events might cause either an immediate shift in approval or a change in its subsequent path, but any given event need not produce such a change. For example, incumbents may adapt before voters act, and reviewers may respond once a successor becomes foreseeable. In Cook County, Foxx entered the race by the late summer of 2015, the McDonald video was released that November, she won the binding Democratic primary in March 2016, and she took office in December. Here, I estimate a series of specifications that assess whether any of these events is associated with a detectable break in a change that may already have been under way.

The specification I estimate involves setting $\nu_a = 0$ for a candidate month c and replacing

the smooth component in Equation (11) with a segmented linear path:

$$\begin{aligned}\mathbb{E}[\text{approve}_{it} \mid a(i), t, x_{it}] = \beta_0 + \beta_1 T_{tc} + \delta_c H_{tc} + \kappa_c T_{tc}^{+} \\ + \gamma\, \text{policy}_{it} + x_{it}'\beta + \widetilde{\alpha}_{a(i)},\end{aligned} \tag{14}$$

where T_{tc} is the number of months from candidate month c, $H_{tc} = \mathbf{1}\{T_{tc} \geq 0\}$, and $T_{tc}^{+} = \max\{T_{tc}, 0\}$. The immediate change in approval that took place at month c is measured by δ_c, while κ_c captures the change in the monthly slope after c. The policy `variable` is an indicator that measures whether a case is a retail theft case occurring after January 2017, as defined above in Equation (12), and to the parameter γ distinguishes that effect from any office-wide change. The vector x_{it} contains linear controls for age and charge severity and fixed effects for race, sex, and broad offense type. I also include reviewer fixed effects. The joint test of a break in the felony approval rate at month c is $H_0 : \delta_c = \kappa_c = 0$.

This specification allows two kinds of inferences. First, it allows the data to identify the month at which a break in approval rates took place by netting the estimated effects of reviewer identity, the retail-theft policy, and observed case mix and then calculating the adjusted approval rate for each month. To do so, I estimate the segmented trend in approval at every month, weighting each monthly observation by its number of cases, and select the month that minimizes the residual sum of squares (Bai 1997).[12] Second, I use the model to estimate the effects of three specific events: the release of the McDonald video, the Democratic primary, and Foxx's inauguration. For each, the candidate month c is the first full month after the event: December 2015 for the video, April 2016 for the primary, and December 2016 for the inauguration (Foxx took office on December 1).[13]

The estimates indicate that the break in felony approval rates took place in September

[12] I estimate standard errors by bootstrapping the fitted model's monthly residuals in six-month blocks and repeating the break-date search for 1,000 resampled series.

[13] I also perform a grid search across all dates as a placebo test and consider shorter estimation windows, which I report in the appendix.

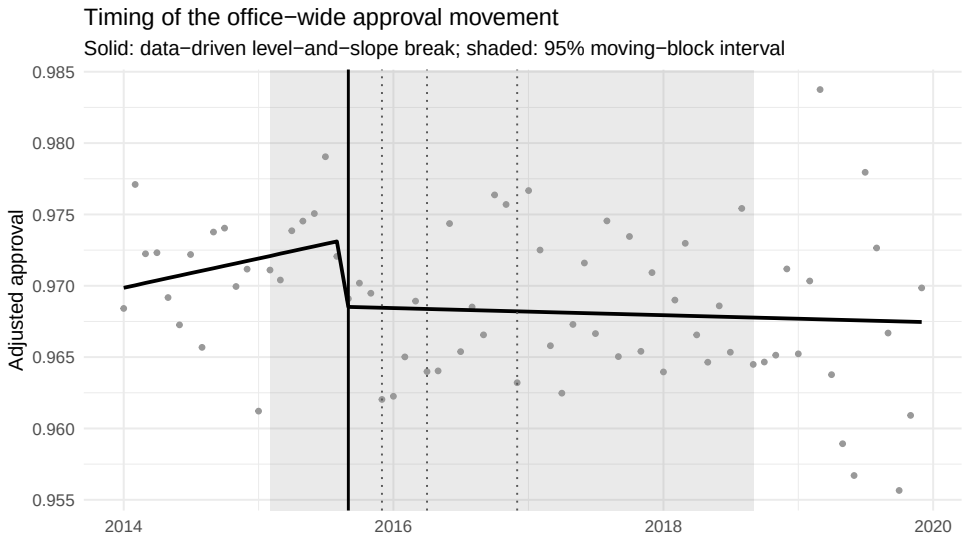


Figure 5: Timing of the office-wide approval movement. The solid line shows the least-squares estimate from Equation (14).. The shaded region is its 95 percent moving-block bootstrap interval, and the dotted lines show the McDonald video, Democratic primary, and inauguration.

2015; however, the confidence interval ranges from February 2015 to September 2018 (Figure 5). That is, the change in felony approval began well before the video, primary, and inauguration, even while the confidence interval contains all three. Thus, a data-driven approach to estimating the timing of a break in approval suggests a decline in approval began before Foxx took office, even if it does not appear to be associated with a single political event.

However, the candidate-date tests yield more precise estimates. In the first full month after the McDonald video's release, adjusted approval falls by 0.59 percentage points. The immediate drop, combined with the continued drop over time attributed to that date, means that a year after the video there had been an aggregate 1.03-point drop in felony review approval.[14] Of course, the video affected both the office's reputational environment and the incumbent's electoral position. As the discussion above underscored, these estimates cannot separate those two pressures. At the same time, neither the primary nor the inauguration

[14]The video ranks third among the 48 admissible dates in the placebo grid, and the result persists in twelve-, eighteen-, and twenty-four-month windows.

Table 3: Candidate-date changes in felony-review approval. Entries are percentage-point estimates with 95 percent confidence intervals. Each row is a case-level linear probability model with reviewer fixed effects, case controls, the retail-theft policy indicator, and a full-period segmented linear trend. The annual slope change is twelve times the monthly post-date slope coefficient. The final column jointly tests the level and slope changes. Newey–West standard errors aggregate scores by month and allow six months of serial dependence.

Candidate date	Immediate shift	Annual slope change	Change after one year	Joint p
McDonald video	−0.59 [−0.94, −0.24]	−0.44 [−1.04, 0.16]	−1.03 [−1.70, −0.36]	0.001
Democratic primary	0.16 [−0.40, 0.72]	−0.35 [−1.08, 0.39]	−0.18 [−1.11, 0.74]	0.553
Inauguration	0.14 [−0.50, 0.78]	−0.68 [−1.38, 0.02]	−0.54 [−1.41, 0.33]	0.162

is associated with a comparable break. The primary is associated with a +0.16 percentage point shift, with an estimated total one-year change of −0.18 points, but this change is not statistically distinguishable from a null effect. Similarly, the immediate change associated with the inauguration is +0.14 points, and the change after one year is −0.54, which also is not distinguishable from a null (Table 3). For both the primary and the inauguration, equivalence tests rule out immediate shifts of at least one percentage point. Taken together, these estimates are more consistent with gradual change after Foxx's inauguration than a large change in response to the election.[15]

Together, these findings suggest that the change in felony review had begun before both the primary and inauguration. There is also no evidence that the trend changed with either event. That does not rule out electoral influence, of course, as the incumbent could have adapted in response to electoral pressure or because Foxx's candidacy made it foreseeable there would be a new state's attorney in the near future. In short, the McDonald video had implications for both the office's immediate reputational environment and the campaign. But, because we do not have exogenous sources of those separate forces, the results only allow us to infer that change in felony review predated Foxx's arrival in office, rather than

[15]The full-period design has 80 percent power to detect an immediate shift of about 0.91 points at inauguration, but a gradual one-year change would have to reach about 1.24 points to be detected with the same power. Eighteen- and twenty-four-month windows also show no joint break at either date. Twelve-month windows fluctuate because they contain adjacent events; Appendix F reports the full sensitivity analysis.

attributing causality to one or the other mechanism.

6 Discussion and Conclusion

Taken together, the results here emphasize and clarify often overlooked theoretical aspects of standard political science accounts of the politics of bureaucratic control. In particular, they underscore that a bureaucracy can change before political authority formally changes hands. Felony-review approval in Cook County declined mainly because the reviewers already working in the office became less likely to approve charges. The data here indicate that changes in personnel and the cases they handle play a much smaller role in driving the decline in felony approval. Of course, Foxx's retail-theft rule produced the sharp response expected for the cases it directly targeted, but the broader decrease in felony approval was already under way when she won the primary and took office. Foxx therefore inherited an office that had begun to change through the decisions of incumbent reviewers.

This finding has deep implications for theories of political control of bureaucratic decision-making. Standard theories of agency and oversight make evident that bureaucrats have strong incentives to respond to their principals' preferences. However, in contrast to most empirical and theoretical analyses of principal-agent control of a bureaucracy, these findings also highlight the ways in which the same logic applies when career officials anticipate the principal they expect to serve. There is, for example, no clear evidence that hierarchical pressure led to broad, office-wide change in Cook County. At the same time, the evidence suggests that the absence of turnover is not necessarily evidence that the office remained unchanged. Instead, the findings show that bureaucratic change can appear in the discretion exercised by incumbent officials even when its political source cannot be identified.

The distinction between personnel authority and personnel replacement also bears on debates over whether removal authority secures democratic control of administration. Observed replacement is an incomplete measure of control because the prospect of personnel

action may change conduct among those who remain. This implication extends to virtually any bureaucracy in which political leaders who come to office and oversee officials who are already in place and exercise substantial discretion. Indeed, the central contribution here the demonstration of how a formal transition in leadership may appear as a lagging indicator of bureaucratic change. The mechanism I highlight is adaptation by an incumbent in advance of an election, just as career officials may respond to an anticipated successor. When external events that shape the reputational environment surrounding the bureaucracy occur, both pressures can induce bureaucratic change, independent of and before an actual transfer of power to a new political principal. In this sense, office-taking is a part of a broader political process of change, rather than the source of it.

This creates an identification problem. Analyses that treat political transitions as empirical indicators of change may describe the difference associated with a new political period, but they do not necessarily isolate the incoming leader's contribution. The pre-inauguration period may already contain adaptation by the incumbent or anticipation among career officials. This is a general lesson for any study of bureaucratic control, but it has particular bite in the politics of criminal justice, an area in which bureaucratic decisions have exceptionally high stakes. In Cook County, the release of the McDonald video changed the CCSAO's reputational environment while simultaneously weakening the incumbent's electoral position. That means that I cannot empirically distinguish a reputational response from reviewers' adapting their behavior in anticipation of Foxx's election. That does not mean that Foxx was irrelevant, of course. However, the evidence does show that neither the primary nor her inauguration identifies the onset of the office-wide change or her contribution to that change. This observation, in turn, has direct implications for empirical studies, including the growing literature evaluating criminal justice and prosecutorial reform. Studies centered on the election of a new prosecutor variously estimate how imprisonment or charging changed across political regimes (Agan et al. 2025, Gottlieb et al. 2025). That is a consequential estimand, but it differs from the effect attributable to the incoming prosecutor after accounting

for adaptation that preceded her. Prosecutorial outcomes also arise from a selected enforcement pipeline. Police determine which cases reach the prosecutor, and there are reasons to suspect that changes in prosecutorial policies may shape the set of cases that police present to a prosecutor in the first instance (Clark Forthcoming, Goldrosen 2026). Thus, combining data on policing and arrests with prosecutorial charging decisions can help separate changes in the cases presented to a prosecutor from changes in how prosecutors use their discretion (Knox, Lowe and Mummolo 2020, Ouss and Rappaport 2020). This is an important avenue for continuing research on the politics of justice and law enforcement, beyond bureaucratic control.

References

Agan, Amanda, Jennifer L. Doleac and Anna Harvey. 2023. "Misdemeanor Prosecution." *Quarterly Journal of Economics* 138(3).

Agan, Amanda Y., Jennifer L. Doleac, Anna Harvey, Anna Kyriazis and Lauren Schechter. 2025. Prosecutorial Reform and Local Crime Rates. NBER Working Paper 34364 National Bureau of Economic Research.

Bai, Jushan. 1997. "Estimation of a Change Point in Multiple Regression Models." *The Review of Economics and Statistics* 79(4):551–563.

Birkland, Thomas A. 1998. "Focusing Events, Mobilization, and Agenda Setting." *Journal of Public Policy* 18(1):53–74.

Bolton, Alexander, John M. de Figueiredo and David E. Lewis. 2021. "Elections, Ideology, and Turnover in the U.S. Federal Government." *Journal of Public Administration Research and Theory* 31(2).

Bowman, Rachel A., Belén Lowrey-Kinberg and Jon B. Gould. 2026. "Progressive Prosecution in Context: Examining the Impact of Prosecutorial Administration Change on

Case-Processing Patterns and Racial Disparities." *Criminology & Public Policy* . Published online.

Brehm, John and Scott Gates. 1997. *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public*. University of Michigan Press.

Busuioc, E Madalina and Martin Lodge. 2016. "The reputational basis of public accountability." *Governance* 29(2):247–263.

Carpenter, Daniel P. 2010. *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton University Press.

Clark, Tom S. Forthcoming. "A Model of Criminal Justice Discretion." *Public Choice* . doi:10.1007/s11127-026-01468-5.

Gailmard, Sean and John W. Patty. 2012. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15:353–377.

Gilad, Sharon, Moshe Maor and Pazit Ben-Nun Bloom. 2015. "Organizational Reputation, the Content of Public Allegations, and Regulatory Communication." *Journal of Public Administration Research and Theory* 25(2):451–478.

Goldrosen, Nicholas. 2026. "The Interorganizational Effects of Reform Prosecution Declination Policies." *Justice Quarterly* . Published online.

Gordon, Sanford C. and Gregory A. Huber. 2002. "Citizen Oversight and the Electoral Incentives of Criminal Prosecutors." *American Journal of Political Science* 46(2).

Gottlieb, Aaron, Matt Epperson, Laura Giugno and Hye-Min Jung. 2025. "What happens to imprisonment rates when a progressive prosecutor is elected: Quasi-experimental evidence from Cook County." *Punishment & Society* 27(2):315–337.

Huber, John D. and Charles R. Shipan. 2002. *Deliberate Discretion? The Institutional Foundations of Bureaucratic Autonomy*. Cambridge University Press.

Imbens, Guido W. and Charles F. Manski. 2004. "Confidence Intervals for Partially Identified Parameters." *Econometrica* 72(6):1845–1857.

Jordan, Andrew. 2024. "Racial Patterns in Approval of Felony Charges." Working paper, Washington University in St. Louis.

Jordan, Andrew. 2025. "Relaxing Charging Restrictions in Narcotics Cases." Working paper, Washington University in St. Louis.

Kingdon, John W. 1984. *Agendas, Alternatives, and Public Policies*. Little, Brown.

Knox, Dean, Will Lowe and Jonathan Mummolo. 2020. "Administrative Records Mask Racially Biased Policing." *American Political Science Review* 114(3).

Lee, David S. 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects." *Review of Economic Studies* 76(3):1071–1102.

Lewis, David E. 2008. *The Politics of Presidential Appointments: Political Control and Bureaucratic Performance*. Princeton University Press.

Lipsky, Michael. 1980. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services*. Russell Sage Foundation.

Manski, Charles F. 1990. "Nonparametric Bounds on Treatment Effects." *American Economic Review* 80(2):319–323.

McCubbins, Mathew D, Roger G Noll and Barry R Weingast. 1987. "Administrative procedures as instruments of political control." *The Journal of Law, Economics, and Organization* 3(2):243–277.

Oberfield, Zachary W. 2014. *Becoming Bureaucrats: Socialization at the Front Lines of Government Service*. University of Pennsylvania Press.

Ouss, Aurélie and John Rappaport. 2020. "Is Police Behavior Getting Worse? Data Selection and the Measurement of Policing Harms." *Journal of Legal Studies* 49(1).

Petersen, Nick, Ojmarrh Mitchell and Shi Yan. 2024. "Do progressive prosecutors increase crime? A quasi-experimental analysis of crime rates in the 100 largest counties, 2000–2020." *Criminology & public policy* 23(2):459–490.

Richman, Daniel. 2003. "Prosecutors and Their Agents, Agents and Their Prosecutors." *Columbia Law Review* 103(4).

Sasso, Greg. 2020. "Delegation and Political Turnover." *Journal of Theoretical Politics* 32(2).

Shjarback, John A, David C Pyrooz, Scott E Wolfe and Scott H Decker. 2017. "De-policing and crime in the wake of Ferguson: Racialized changes in the quantity and quality of policing among Missouri police departments." *Journal of criminal justice* 50:42–52.

Sklansky, David Alan. 2018. "The Problems with Prosecutors." *Annual Review of Criminology* 1.

Weingast, Barry R. and Mark J. Moran. 1983. "Bureaucratic Discretion or Congressional Control? Regulatory Policymaking by the Federal Trade Commission." *Journal of Political Economy* 91(5).

Wood, B Dan and Richard W Waterman. 1991. "The dynamics of political control of the bureaucracy." *American Political Science Review* 85(3):801–828.

Online Appendix

Political Control and the Sources of Bureaucratic Change

Appendix Contents

A The Charging-Standard Model A-2

B Reviewer Assignment and Balance A-5

C Baseline Specification and Reviewer Effects A-8

D Varying the Composition of Reviewers A-9

E Pooled Retail-Theft Policy Specification A-10

F Timing Tests A-12

G Does the Changing Set of Cases Drive the Results? A-15

A The Charging-Standard Model

Here, I develop the incumbent's dynamic problem and derive the reviewer's optimal review decision.

Reviewer choice. Consider an implemented standard, s_t. Reviewer a chooses an approval strategy to minimize the loss in Equation (3). The first-order condition is

$$(\tau - \alpha_a) + \kappa(\tau - s_t) = 0.$$

Writing $w = \kappa/(1+\kappa)$ gives

$$\tau^0_{at}(s_t) = (1-w)\alpha_a + w s_t,$$

which is reported in the paper as Equation (6). To see this best response is unique, note that the second derivative is $2(1+\kappa) > 0$.

Source-specific heterogeneity. Consider an extension in which reviewer a has some predetermined degree of career exposure to the principal, denoted z_a and allow for the principal to adjust the office standard for each reviewer-case pair in a way that is related to z_a. Let $h(o_i) \in \mathbb{R}$ denote the exposure of offense o_i to reputational scrutiny. Further, define both measures as deviations from their case-weighted means, normalized so that average exposure is zero. With this extension, the principal's choice of an office standard specific to a on case i becomes

$$\widetilde{s}^*_{ait} = \widetilde{s}^*_t + \lambda_B z_a B_t + \lambda_E h(o_i) E_t, \qquad \lambda_B, \lambda_E \geq 0, \tag{15}$$

where λ_B and λ_E are parameters. Substituting this office standard into the reviewer's best response gives

$$\tau^*_{ait} = (1-w)\alpha_a + w\widetilde{s}^*_{ait}, \tag{16}$$

and the cross-partial derivatives are given by

$$\frac{\partial^2 \tau^*_{ait}}{\partial B_t \, \partial z_a} = w\lambda_B, \qquad \frac{\partial^2 \tau^*_{ait}}{\partial E_t \, \partial h(o_i)} = w\lambda_E.$$

These results imply we can distinguish hierarchical and environmental responses only if z_a modifies the response to hierarchy alone and $h(o_i)$ modifies the response to reputational pressure alone. Unfortunately, the extant data do not contain measures that satisfy those restrictions. Centering the exposure variables on zero implies we can replace $g(t) = -\bar{f}ws^*_t$ with $g(t) = -\bar{f}w\widetilde{s}^*_t$ in Equation (9) but does not otherwise change the aggregate decomposition.

Thus, in this model, the reviewer's approval choice is given by

$$Y_{ait} = \mathbf{1}\{v_i > \tau^*_{ait}\}\,\mathbf{1}\{D_{it} = 0\}. \tag{17}$$

Dynamic office choice and equilibrium. Let s denote the existing standard, and let Z denote the political state. For an incumbent with intrinsic standard $\bar{s}$, the continuation value is given by

$$\mathcal{C}_{\bar{s}}(s'; s, Z) = \mathbb{E}\left[\mathbf{1}\{\text{retained}\}V_{\bar{s}}(s', Z') + \mathbf{1}\{\text{replaced}\}\overline{V}_{\bar{s}} \mid s, Z, s'\right].$$

We can write Equation (5) as

$$V_{\bar{s}}(s, Z) = \max_{s'} \left\{-\ell_{\bar{s}}(s'; s, E) + \delta\mathcal{C}_{\bar{s}}(s'; s, Z)\right\}.$$

The first order condition (for an interior choice) is

$$2\big(s' - \hat{s}(E)\big) + 2\zeta(s' - s) = \delta\frac{\partial\mathcal{C}_{\bar{s}}(s'; s, Z)}{\partial s'}. \tag{18}$$

The left side contains the incumbent's current ideal point and the cost of changing the inherited standard. The right side is the effect of a change in today's standard on her political future.

Denote by $p_{\bar{s}}(s', Z)$ denote the probability that the incumbent is retained at an election. The continuation value includes

$$p_{\bar{s}}(s', Z)\,\mathbb{E}[V_{\bar{s}}(s', Z') \mid \text{retained}] + \big(1 - p_{\bar{s}}(s', Z)\big)\overline{V}_{\bar{s}},$$

and its derivative contains $\frac{\partial p_{\bar{s}}(s', Z)}{\partial s'}(\mathbb{E}[V_{\bar{s}}] - \overline{V}_{\bar{s}})$. A standard that improves the incumbent's reelection probability changes her standard choice while in office because she values remaining in office. That logic implies even earlier changes can affect her reelection to the extent they influence the political state at the time of an election. And, a new successor makes her initial choice having inherited the standard chosen by the prior incumbent.

Proposition 1 *Suppose $\delta \in (0, 1)$, the state and standard spaces are compact, all standards are feasible, payoffs are bounded and continuous, and the transition kernel is weakly continuous. A Markov-perfect Stackelberg equilibrium exists. If the incumbent's Bellman objective has a unique maximizer, she uses the policy $s_t^* = \sigma_{\bar{s}_t}(s_{t-1}^*, Z_t)$. Each reviewer $a \in \mathcal{R}_t$ chooses $\tau_{at}^* = (1 - w)\alpha_a + w s_t^*$.*

Proof. Substituting the unique reviewer best response from Equation (6) into the incumbent's payoff reduces the leader's problem to a discounted Markov decision problem. Under the stated assumptions, its Bellman operator maps bounded continuous functions into themselves and is a contraction with modulus δ. It therefore has a unique bounded continuous fixed point. Compactness and continuity imply that the maximization in the Bellman equation attains its value and admits a stationary Markov selection. Pairing that policy with each reviewer's unique best response yields a Markov-perfect Stackelberg equilibrium. When the maximizer is unique, the selection is the policy $\sigma_{\bar{s}_t}(s_{t-1}^*, Z_t)$ stated in the proposition. □

Linear-quadratic benchmark. Finally, consider a linear-quadratic benchmark by holding the political regime and ideal point fixed. Suppress electoral transition and denote the

incumbent's value function as $J(s)$. The dynamic problem becomes

$$J(s) = \min_{s'} \left\{ (s' - \hat{s})^2 + \zeta(s' - s)^2 + \delta J(s') \right\}.$$

The solution has $J(s) = P(s - \hat{s})^2 + C$, where the positive value of P solves

$$\delta P^2 + \big(1 + \zeta - \delta\zeta\big)P - \zeta = 0.$$

The policy is

$$\sigma(s) = (1 - \rho)\hat{s} + \rho s, \qquad \rho = \frac{\zeta}{1 + \zeta + \delta P} \in [0, 1).$$

As a result, the previous time's standard partially affects the current standard choice even when the principal's ideal point is constant. When the political state is expected to change, the continuation term in Equation (18) also affects the current standard, which demonstrates that changes to the standard can begin before an anticipated election.

B Reviewer Assignment and Balance

The decomposition of the charging shift into personnel and common-movement components relies on the claim that felony-review cases are assigned to on-call reviewers in a manner close to random. Institutionally, reviewers work fixed twelve-hour shifts and incoming calls are assigned in rotation to whoever is next in the shift's order, so assignment depends on the order of calls and the roster (Section 4; Jordan 2024). Conditional on the active roster and call time, observed case attributes should therefore be unrelated to reviewer identity. I do not observe the roster or call time, so the tests below use review month as a coarse proxy for the assignment block.

Panel A of Table A1 reports joint F-tests of the reviewer indicators in separate regressions for each observed case attribute. The null of identical conditional means across reviewers is rejected for every attribute. For the attributes other than narcotics, the F-statistics range

from 1.29 to 3.58. The narcotics statistic is 100.90, a qualitatively different association between reviewer identity and offense type. Adding offense-category fixed effects reduces the test statistics for the remaining attributes, but each test still rejects at $p < .001$.

These rejections show the limit of the available balance test. The identifying comparison is among cases arriving under the same active roster at similar times, while month fixed effects compare cases across much broader assignment conditions. The joint tests therefore document detectable differences in observed case mix but do not reconstruct the assignment blocks or establish exact balance. Jordan (2024) documents the rotation and reports complementary balance tests in an overlapping sample.

The narcotics row also reflects a difference between two classifications used in the paper. Table A1 uses the prosecutor's broad offense category for a reviewed case. The selection analysis classifies an arrest as narcotics-only only when every recorded felony arrest charge matches one of the specified narcotics statutes. Mixed-charge arrests can therefore be eligible for review even when the reviewed case receives a narcotics label. The linked data also contain a small number of narcotics-only arrests that appear in the review file; the data cannot distinguish exceptions to the routing rule from inconsistencies between arrest-statute and prosecutor-case coding. In the baseline complete-case sample, narcotics account for only 217 of 93,378 decisions, but they are concentrated among a small number of reviewers. Panel B therefore reestimates the decomposition after excluding every narcotics-category decision and reapplying the minimum-caseload rule. The adjusted decline is 4.00 percentage points, of which 16.9 percent is attributed to personnel, compared with 4.05 points and 18.7 percent in the full sample. The imbalance in this small category does not account for the decomposition.

Because felony review is a rotating junior assignment, the reviewer panel is not a fixed cast, and the within-reviewer results should be read as movement within observed tenures rather than as following a large set of literal stayers across administrations. Table A2 reports the structure of the panel, distinguishing the reviewer count before complete-case restrictions

Table A1: Reviewer assignment and sensitivity to narcotics cases.

Panel A: Joint tests of observed case attributes

	Month fixed effects		Month and offense fixed effects	
Case attribute	F	p	F	p
Charge severity	3.58	$< .001$	2.58	$< .001$
Defendant age	2.27	$< .001$	1.51	$< .001$
Black defendant	3.58	$< .001$	2.95	$< .001$
Male defendant	1.41	$< .001$	1.29	$< .001$
Class M/X or 1 charge	2.95	$< .001$	2.24	$< .001$
Narcotics offense	100.90	$< .001$	–	–

Panel B: Decomposition sensitivity

Sample	Cases	Reviewers	Total	Personnel	Common	Personnel share
All offenses	93,378	321	−4.05	−0.76	−3.29	18.7%
Excluding narcotics	93,161	320	−4.00	−0.68	−3.33	16.9%

Panel A reports the conventional joint F-test and p-value for the null that all reviewer indicators equal zero in a linear regression of the listed attribute. The models use 93,378 cases and 321 reviewers. The reviewer test has 320 numerator degrees of freedom; the residual degrees of freedom are 92,986 with month fixed effects and 92,982 after adding offense-category fixed effects. The offense-conditioned test is not defined for the narcotics indicator because that indicator is a function of offense category. In Panel B, total, personnel, and common changes are in percentage points; the 25-case reviewer threshold is reapplied after excluding narcotics.

from the exact baseline estimation sample. It also reports the typical caseload per reviewer, the median observed tenure, the number of reviewers who handle cases on both sides of the December 2016 inauguration, and the share of all reviewed cases those spanning reviewers handle.

Table A2: Structure of the reviewer panel, 2014–2019.

Quantity	Value
Reviewers observed	550
Reviewers with $\geq$ 25 decisions before complete-case restrictions	324
Reviewers in baseline estimation sample	321
Median cases per reviewer	56
Median observed tenure (months)	19
Reviewers observed on both sides of the inauguration	224
Share of all reviewed cases handled by spanning reviewers	0.65
Share of cases in 2016–2017 handled by spanning reviewers	0.89

C Baseline Specification and Reviewer Effects

Table A3 reports the primary estimated coefficients from Equation (12), with standard errors clustered by review month. (The model also includes reviewer fixed effects and the twelve-degree-of-freedom natural spline in review month.)

In Figure A1, I plot the reviewer unweighted fixed-effect estimates. (I centered them on their case-weighted mean, because the particular values depend on which reviewer constitutes the omitted case.)

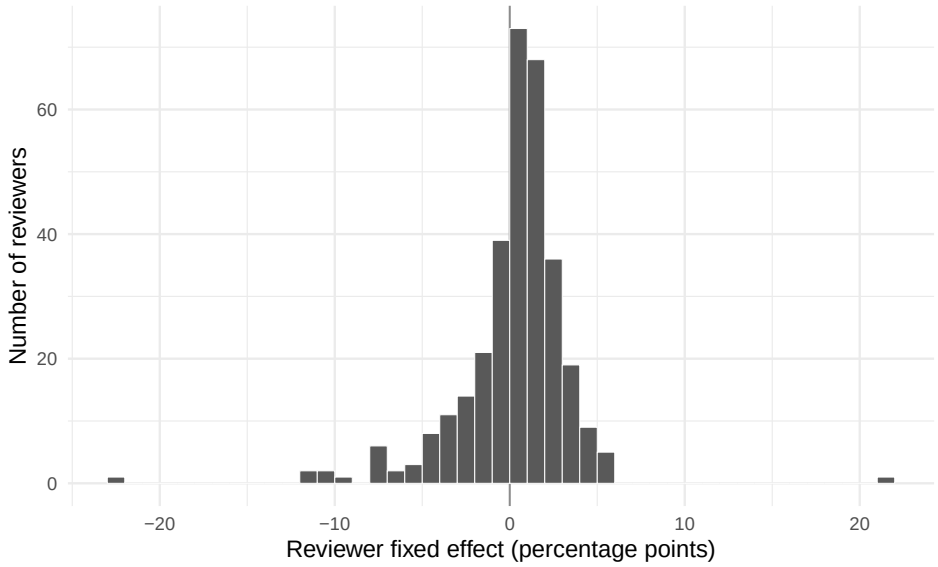


Figure A1: Distribution of estimated reviewer fixed effects from Equation (12). Effects are centered at their case-weighted mean and expressed in percentage points. The sample contains 321 reviewers with at least twenty-five cases.

Table A3: Baseline linear probability model of felony-review approval.

Variable	Estimate	Std. error	95% CI
Intercept	95.839	3.184	[89.597, 102.080]
Retail theft × January 2017 onward	−4.957	0.587	[−6.107, −3.808]
Defendant age (years)	0.014	0.005	[0.004, 0.024]
Charge severity	0.222	0.052	[0.120, 0.324]
Hispanic defendant	0.385	0.144	[0.103, 0.667]
Other-race defendant	−0.508	0.821	[−2.118, 1.101]
White defendant	−0.157	0.184	[−0.517, 0.203]
Male defendant	0.283	0.178	[−0.066, 0.632]
Sex unknown	−0.470	3.677	[−7.677, 6.737]
Narcotics offense	−19.729	3.710	[−27.000, −12.459]
Non-violent offense	0.055	0.192	[−0.322, 0.431]
Sex offense	−2.052	0.535	[−3.100, −1.005]
Violent offense	−0.290	0.229	[−0.740, 0.159]
Observations			93,378
Reviewers			321
Review months			72
Reviewer fixed effects			Yes
Natural spline degrees of freedom			12
R^2			0.034

Coefficients and standard errors are reported in percentage points. Standard errors cluster by review month. The omitted categories are Black defendants, female defendants, and gun offenses. The model includes reviewer fixed effects and a twelve-degree-of-freedom natural spline in review month.

D Varying the Composition of Reviewers

A simple alternative to the paper's main specification is to split the pre-to-post change in approval into a between-reviewer component and a within-reviewer component by using the the changing share of cases each reviewer handles and changes in each reviewer's own rate. In the context of felony review, this results in an estimate that roughly two-thirds of the decline in felony review approval is due to composition. However, because there is change in who serves as a felony reviewer, the estimate is mechanically inflated, so it should not be interpreted as evidence that replacement explains two-thirds of the decline.

The challenge is that felony review is a junior assignment, and assistant state's attorneys tend not to stay in that role for very long. As a result, we only observe most reviewers on one side of the December 2016 inauguration (Table A2). We cannot estimate a pre-transition

approval rate for any reviewer that does not appear in the sample before December 2016, and we must impute one, and the natural imputation is the pre-period average. This group comprises a very large share of the reviewers we observe after the inauguration (Table A2). For these reviewers, we can only assume their approval rates are comparable to those before the transition, including the over-time declining trend in approval. However, because they have no measured pre-period baseline, any temporal trend associated with such reviewers appears as a composition effect, rather than as within-reviewer change. In other words, we can only estimate within-reviewer contributions to the declining approval rate for reviewers who appear on both sides of the transition, and so this specification mechanically over-estimates the role composition plays in explaining the declining approval rate, whether or not the incoming reviewers differ from the outgoing ones.

The main specification in the paper avoids this imputation by estimating fixed effects for each reviewer after removing the common change over time and case mix and weighting those effects by the reviewer's actual case shares in 2014 and 2019. The estimated effect for a reviewer who appears on only one side of the transition is determined only by his own decisions net of the temporal effects. As reported in the paper, this approach estimates that 18.7% of the decline in approval is attributable to reviewer composition, compared with 81.3% for a trend in the office standard. These findings are stable across a variety of specifications, including variation in the degrees of freedom in the time trend and varying the minimum number of cases for a reviewer to be included in the sample from 10 to 100.

E Pooled Retail-Theft Policy Specification

Here I analyze the sample of cases that are complete cases—i.e., have non-missing values for the review outcome, review date, reviewer, defendant age, charge severity, race, sex, and offense category. This sample overs felony-review decisions from January 1, 2014, through December 31, 2019. I limit the analysis to reviewers who handled at least 25 cases. In the

main specification further restricts the sample to cases classified as retail theft or ordinary theft. The result is a sample of 14,125 decisions by 312 reviewers across 72 month-years.

Let R_i indicate retail theft and let $P_{it} = \mathbf{1}\{d_i \geq \text{December 15, 2016}\}$, where d_i is the review date. I estimate

$$\mathbb{E}[\text{approve}_{it} \mid a(i), t, x_{it}] = \mu_t + \rho R_i + \lambda P_{it} + \gamma(R_i \times P_{it}) + x'_{it}\beta + \widetilde{\alpha}_{a(i)}, \tag{19}$$

where μ_t is a month-year fixed effect and $\widetilde{\alpha}_{a(i)}$ is a reviewer fixed effect. The vector x_{it} includes the linear measures of the defendant's age and the charge severity, as well as indicators for race and sex. The retail-theft indicator R_i captures the baseline difference between the two offenses. Month-year effects capture changes across both categories over time. Because the policy begins halfway through December 2016, P_{it} also varies within that month; includong it captures any change after December 15 shared by retail and ordinary theft. The interaction coefficient γ is the differential post-policy change in retail-theft approval. I cluster standard errors by month-year.

The main specification uses the policy's effective date. I estimate three alternatives. First, I begin the post-policy period in January 2017, the first full month after the rule. As a result, that indicator is perfectly collinear with the month-year effects, so in that specification I only include its interaction with retail theft. Second, I expand the set of comparison offenses to include a broader property-offense comparison group: theft, burglary, residential burglary, possession of a stolen motor vehicle, forgery, identity theft, credit-card cases, criminal damage to property, fraud, deceptive practice, theft by deception, and aggravated identity theft. Third, I further expand the sample to include all offense categories except driving with a suspended or revoked license, which was subject to a separate policy during the study period.

The result of this analysis are summarized in Table A4. Compared to ordinary theft, retail-theft approval falls by approximately 3.5 percentage points after December 15. If we instead measure the post period beginning in January 2017, the estimate is a roughly

3.2-point decline. Comparing retail theft to other property crime offenses, the estimate is about 4 points; comparing to all other offenses (other than driving with a suspended license) is about 5 points. Importantly, though, there is measurement error in which offenses are covered by the policy, because the data do not include ever relevant factor for whether an offense is covered (e.g., prior convictions, exact amount stolen, etc.)

Table A4: Pooled estimates of the retail-theft policy response.

Comparison	Difference (p.p.)	95% CI	Cases
Ordinary theft; Dec. 15 cutoff	−3.56	[−5.77, −1.34]	14,125
Ordinary theft; Jan. 2017 cutoff	−3.26	[−5.53, −1.00]	14,125
Property offenses; Dec. 15 cutoff	−4.08	[−5.38, −2.77]	30,923
All offenses except DWLS; Dec. 15 cutoff	−5.04	[−6.22, −3.85]	88,997

Each row reports the coefficient on retail theft interacted with the indicated post-policy period. All models include reviewer and calendar-month fixed effects, linear controls for age and charge severity, and fixed effects for race and sex. Models using the December 15 cutoff also include the post-date main effect. Standard errors cluster by calendar month. The estimation sample and comparison groups are described in the text.

F Timing Tests

In the main specification, I use the full period to account for gradual adjustments that take place between the various events that occurred near to each other, estimating both immediate and year-long trends. In panel A of Table A5, I allow the window around each event to narrow. By limiting the panel to only 12 months, I reject the null of no change for each of the three events—the McDonald video, the primary, and the inauguration. However, we cannot use these tests to distinguish tghe three events: each twenty-five-month estimation window contains the video, the primary, and the inauguration, while the primary and inauguration windows also contain the retail-theft policy change. Panel B removes all retail-theft decisions and again rejects at all three dates. The narrow-window results therefore are not driven solely by the retail-theft rule; they show that a single unbroken linear path fits the densely overlapping period poorly. With eighteen-month, twenty-four-month, or full-period

windows, the video result persists, while neither the primary nor the inauguration produces a detectable joint change.

In Panel C, I show the results of a power analysis reporting minimum detectable effects with 80 percent power, in percentage points. The full-period specification can detect immediate shifts of 0.50 points at the video, 0.80 at the primary, and 0.91 at inauguration. The data are less informative about gradual changes, though Here, the minimum detectable one-year change is 1.32 points at the primary and 1.24 points at inauguration. Ninety-percent equivalence intervals fall within a one-point bound for both immediate electoral-date shifts. The annual inauguration slope and its combined one-year effect do not meet that equivalence standard.

Panel D reports the three highest joint statistics in the full-period placebo grid. September 2015 ranks first, November 2015 ranks second, and the December 2015 video month ranks third. The concentration of the highest-ranked dates in late 2015 supports locating the change in that period, but it does not identify the video-release month uniquely or establish that the video caused the change.

Table A5: Sensitivity, detection limits, and placebo dates for candidate-date timing tests.

	Estimation window around candidate date			
Candidate date	12 months	18 months	24 months	Full period
Panel A: Joint p-value for level and slope change				
McDonald video	$< .001$	$< .001$	$< .001$	0.001
Democratic primary	0.001	0.122	0.700	0.553
Inauguration	$< .001$	0.292	0.203	0.162

Panel B: Twelve-month-window sensitivity to excluding retail theft		
Candidate date	All cases	Excluding retail theft
McDonald video	$< .001$	$< .001$
Democratic primary	0.001	0.005
Inauguration	$< .001$	$< .001$

Panel C: Full-period detection limits and placebo rank				
Candidate date	Level MDE	Annual slope MDE	One-year MDE	Placebo rank
McDonald video	0.50	0.86	0.96	3/48
Democratic primary	0.80	1.05	1.32	47/48
Inauguration	0.91	1.00	1.24	25/48

Panel D: Highest-ranked full-period dates				
Rank	Date	Joint statistic	Joint p	Relationship to video
1	September 2015	17.04	$< .001$	3 months before video
2	November 2015	14.17	$< .001$	1 month before video
3	December 2015	13.68	0.001	Video month

Panel A reports the joint p-value for the immediate level and post-date slope terms across estimation windows. Panel B repeats the twelve-month-window tests after removing all retail-theft decisions. Panel C reports 80-percent minimum detectable effects in percentage points and the descending rank of each candidate date's joint statistic among 48 admissible dates. Panel D reports the three highest-ranked dates in that grid. All specifications are case-level linear probability models with reviewer fixed effects and case controls; specifications retaining retail-theft decisions also include the retail-theft policy indicator. Newey–West standard errors aggregate scores by month and use a six-month lag.

G Does the Changing Set of Cases Drive the Results?

In this appendix, I address whether the decline in approval rates simply reflects changing arrest patterns, so that the decline represents review rates in qualitatively different cases. Here, I analyze only the sample of felony review cases that can be linked to arrest data from the Chicago Police Department, rather than the full set of county-wide cases shown in Figure 2. Chicago felony arrests fell from roughly 25,000 in 2014 to roughly 18,000 in 2017, rising afterwards to roughly 22,000 in 2019. The decline most mostly attributable to narcotics cases, which bypass felony review. Among felony-review eligible arrests, the share reaching review fell from about 0.97 in 2014 to about 0.90 in 2019 (Figure A3; Table A6). Moreover, the composition of offenses in these arrests and the distribution of arrestee racial identities change. As a result, the empirical patterns in charging involve changes in both how reviewers exercise disctetion and the set of cases selected into review by police arresting patterns.

We find a particularly concentrated version of this dynamic among the cases covered by the retail-theft policy. Retail theft accounts for about 15% of the cases presented for felony review before the policy was adopted but only about 5% thereafter (Figure A2). That drop is immediate and persistent after the policy adoption. Thus, any change in approval rates after the policy must be interpreted as conditioning on a case pool that changed.

To bound the within-pool change, I estimate trimming bounds (Lee 2009). These bounds impose a monotonicity assumption. Specifically, the assumption requires that when the selection rate falls from s_0 to $s_1 < s_0$, the resulting set of cases is a subset of the cases the looser regime would have reviewed. In other words, the monotonicity assumption is that heightened selection leads to an analysis only of cases that would have been reviewed under either standard. Formally, a looser review standard includes the tighter standard as well as a marginal fraction $q = (s_0 - s_1)/s_0$. However, the data do not inform where those marginal cases are located in the approval distribution. Thus, if we remove the bottom q of the looser period's distribution, we estimate the highest approval rate; deleting the top q yields the

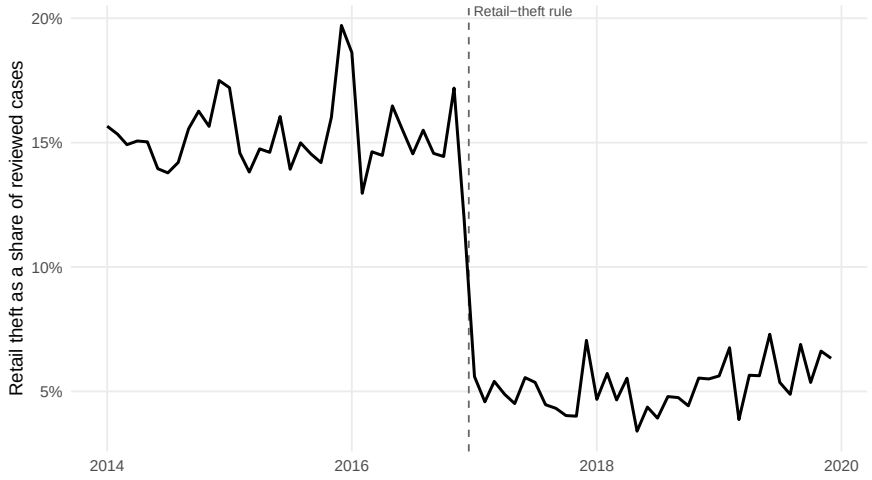


Figure A2: Retail theft as a share of the felony-review docket. The share falls from about 0.15 to about 0.05 when the December 2016 declination policy takes effect and remains near that level through 2019.

lowest. The result is an interval $\Delta_t^L \leq \Delta_t \leq \Delta_t^U$ giving the worst-case with respect to how the marginal cases sort on the outcome. I also estimate results from tightening these bounds via by the same trimming procedure within race-by-offense-class strata.

The bounding exercise indicates that selection alone cannot account for the decline in approval rates. The worst-case interval for the within-pool change lies below zero for every year from 2016 forward, and trimming the bounds by race-by-offense-class strata results in even tighter bounds. Selection also moves in the opposite direction from the decline in approval, as entry into the set of cases reviewed becomes more restrictive over time. If the cases that are still reviewed under a stricter approval standard were those that were inherently stronger, then the decline in approval will be an underestimate of the change in the office standard.

Three further types of analysis corroborate this interpretation. First, reweighting each year's reviewed pool to the 2014 distribution of race and felony class does not affect the estimates in a meaningful way, leading to a less than 1 percentage-point change in the estimates. Second, if I include the unlinked data and code them as being approved in the initial year and declined in later years, increases then magnotude of the decline, indicating

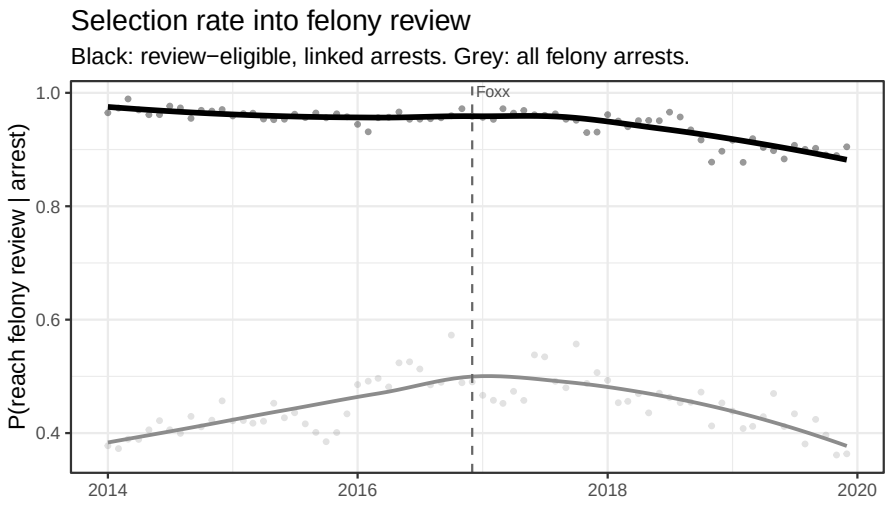


Figure A3: Selection into felony review. Black: among review-eligible arrests, the share reaching felony review declines from about 0.97 in 2014 to about 0.90 in 2019. Gray: all felony arrests, where the level is dominated by narcotics-only arrests that bypass review by rule and by unlinked bookings. The pool entering the funnel is endogenous to the enforcement environment.

the change is not due to unmatched arrests. Third, the estimates hold even if I drop last quarter of 2019.

Linkage diagnostics. I link arrests to felony review outcomes via an arrest-information file obtained through a Freedom of Information Act request. The file contains a cross-walk to link arrest booking numbers with anonymized case-participant identifiers in the State's Attorney's data (Section 4). The match rate averages about 95% through 2018 but declines over time, starting in the latter half of that year, especially in the final quarter of 2019 (Figure A5). Because the selection rate is computed conditional on a successful link, that cannot enter the estimates mechanically. The adverse-coding and late-2019 checks described above demonstrate that show that this decline in matching does not produce the decline in approval rates. The arrests that never link include both genuine non-prosecution and coding errors in the booking numbers.

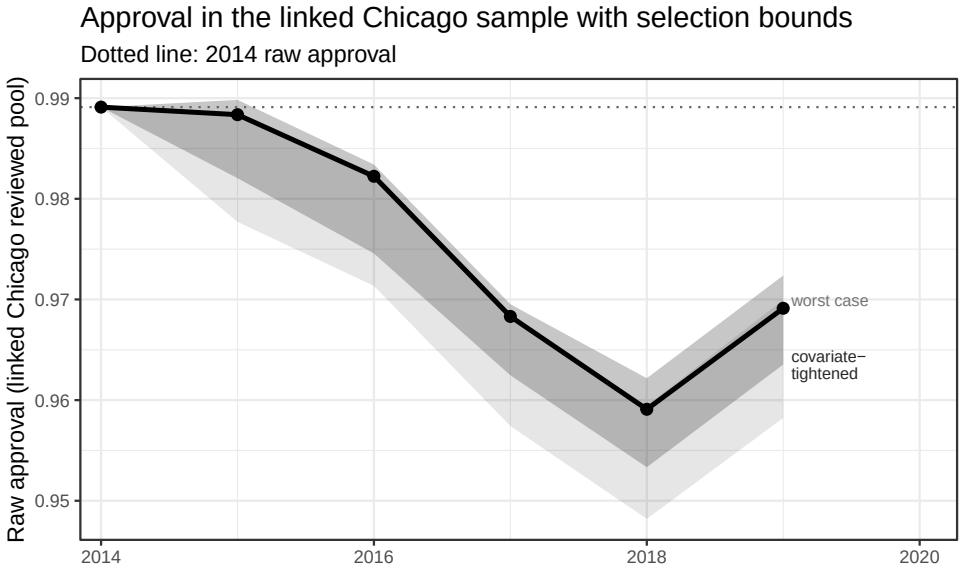


Figure A4: Trimming bounds on raw approval in the linked Chicago reviewed pool, computed from the review-eligible arrest universe. Even the worst-case bounds (light) place the 2016–2019 changes below zero; the covariate-tightened bounds (dark) narrow them further. Both sets impose monotone selection. This figure is not an estimate of the adjusted countywide calendar component $g(t)$.

Table A6: Selection into review and bounds on the within-review change in approval, relative to 2014.

Year	Selection rate S_t	Raw Δ approval	Worst-case [lo, hi]	Covariate-tightened	Imbens–Manski 95% CI
2015	0.959	−0.001	[−0.011, −0.001]	[−0.007, +0.001]	[−0.015, +0.002]
2016	0.955	−0.007	[−0.018, −0.007]	[−0.015, −0.006]	[−0.020, −0.004]
2017	0.956	−0.021	[−0.032, −0.021]	[−0.027, −0.020]	[−0.035, −0.017]
2018	0.938	−0.030	[−0.041, −0.030]	[−0.036, −0.027]	[−0.044, −0.026]
2019	0.900	−0.020	[−0.031, −0.019]	[−0.026, −0.017]	[−0.034, −0.016]

S_t is the annual share of review-eligible felony arrests (narcotics-only arrests, which bypass review by rule, are excluded) that reach felony review, conditioned on a successful arrest-to-case link; the 2014 baseline is 0.969. *Because the selected share is high, even the worst-case bounds sign the decline from 2016 on. Both bound sets impose monotone selection; the worst-case columns impose nothing else. The covariate-tightened bounds trim within race × offense-class strata. The final column reports the Imbens–Manski 95 percent confidence interval for the worst-case bound. From 2014 to 2019 the selection rate fell by* 0.069 *while approval fell by* 0.020*, so positive selection implies the observed decline understates the within-pool shift.*

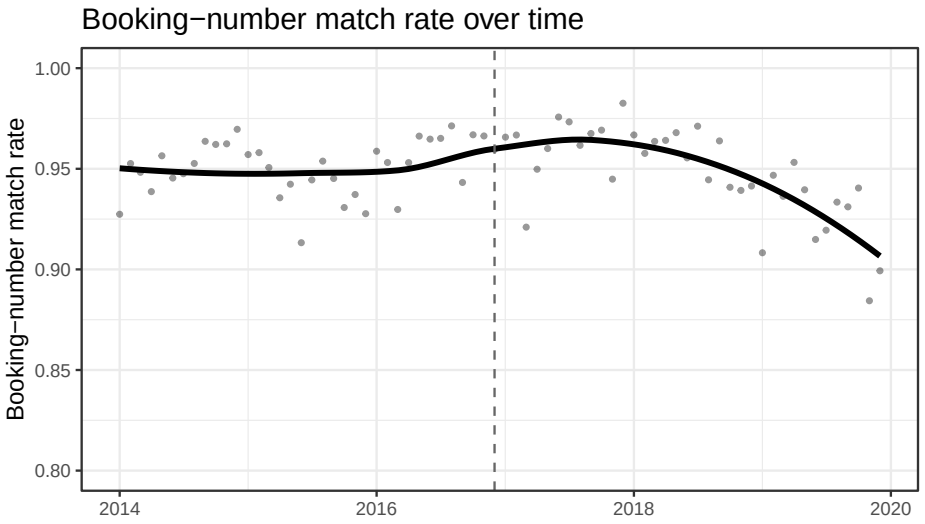


Figure A5: Booking-number match rate by month. The rate averages about 95 percent through 2018, begins drifting downward in the second half of 2018, and falls further in 2019. The selection-rate analysis conditions on a successful link, and the text reports sensitivity checks for the remaining coverage change.