

An Enduring Interruption: How “Zero Tolerance” Reshaped Prosecutorial Discretion at the U.S.–Mexico Border

Matthew Blanton

Department of Sociology and Population Research Center
The University of Texas at Austin

Please direct all correspondence to:

Matthew Blanton
305 E 23 St, A1700, RLP 3.306
Austin TX 78712-1086
mbanton@utexas.edu

Acknowledgements: This research was supported by grants P2CHD042849 and T32HD007081, awarded to the Population Research Center at The University of Texas at Austin by the Eunice Kennedy Shriver National Institute of Child Health and Human Development. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health. I thank Josiah Davis for his instructive feedback on this study's development and design.

An Enduring Interruption: How “Zero Tolerance” Reshaped Prosecutorial Discretion at the U.S.–Mexico Border

Abstract

In April 2018, the U.S. Department of Justice announced a “Zero Tolerance” policy directing federal prosecutors to criminally charge every adult apprehended crossing the U.S.–Mexico border without authorization. Carrying it out required Border Patrol to refer every case, ending the agency's practice of selecting whom to send to prosecutors. The policy lasted barely three months, and per-apprehension referral rates, after spiking, largely receded. Did the way Border Patrol decided whom to refer return to normal when the rate did, or did Zero Tolerance reshape that decision? Using FOIA-obtained Customs and Border Protection records on all 1,000,781 Single Adult apprehensions from January 2016 through December 2019, this study fits gradient-boosted classifiers on rolling three-month windows and decomposes them with SHAP values to track the referral decision before, during, and after the directive. In this analysis, I make three contributions. Methodologically, I offer a rolling-window design that estimates how a discretionary decision rule shifts over time, capturing change a single regression would average away. Substantively, I find that under Zero Tolerance prosecutorial discretion was never eliminated—referrals peaked at only 61%. However, the selection rule was durably reorganized: region of origin gained weight while criminal history and other case factors lost it. Conceptually, I demonstrate that a brief policy interruption can leave a lasting structural residue invisible at the level of the rate.

Keywords: zero tolerance; prosecutorial discretion; immigration enforcement; interpretable machine learning; SHAP

Highlights

- Zero Tolerance aimed to end prosecutorial discretion at the U.S.–Mexico border.
- A rolling-window design tracks the referral decision rule month by month.
- Discretion was compressed, not eliminated: referral peaked at only 61%.
- Region of origin gained weight as criminal history and case factors lost it.
- A brief policy left a durable structural residue invisible in referral rates.

Introduction

On April 6, 2018, Attorney General Jeff Sessions directed federal prosecutors along the U.S.–Mexico border to bring criminal charges against *every* adult apprehended entering the country without authorization. Announced as a policy of “zero tolerance” (hereafter Zero Tolerance, or ZT), the directive was explicit about its purpose: to strip discretion out of the prosecution of illegal entry and replace it with a categorical mandate to refer all cases (U.S. Department of Justice, 2018). The directive implemented President Trump’s Executive Order 13767, which directed agencies to “deploy all lawful means to secure the Nation’s Southwest border” and “prevent further illegal immigration into the United States” (Kandel, 2021). In practice, however, the main consequence of Zero Tolerance was the separation of thousands of children from parents sent to criminal court, a practice that generated intense public opposition and pushed the administration toward a *de facto* reversal about eleven weeks later in June 2018 (Eagly, 2020).

The extent to which border enforcement followed through on the mandate is contested. While a Congressional Research Service report claimed that the DOJ prosecuted “100% of adult aliens apprehended” (Kandel, 2021), other analysis shows that in May 2018, less than one third of apprehended adults were prosecuted (TRAC Immigration, 2018). Even less is known about what the policy did to the referral decision that precedes prosecution—the choice, made by Border Patrol, to route an apprehended adult toward criminal charges in the first place. Did this discretion shut down in spring 2018 and then resume once the policy receded? Or did Zero Tolerance change the sorting process by which border crossers were routed toward prosecution or away from it?

Zero Tolerance was a brief episode. The directive operated for roughly eleven weeks, and the referral rate, after spiking, largely receded toward where the pre-policy escalation had left it. Judged by the rate, the episode looks closed. But a rate can return while the rule that produces it does not: Border Patrol could refer cases at the old frequency while selecting them on a new basis, and the two outcomes would look identical in an aggregate series. Did the pre-policy rule return, or did a different rule come to produce a similar rate? The rate itself cannot say. Answering requires looking *inside* the decision.

To explore these questions, I use interpretable machine learning to trace, at the level of the individual referral decision, how the internal structure of border prosecution changed before, during, and after Zero Tolerance. I use individual-level apprehension records obtained from U.S. Customs and Border Protection (CBP) through the Freedom of Information Act: 1,000,781 Single Adult apprehensions by Border Patrol, the CBP component that polices between ports of entry, spanning January 2016 through December 2019. The decision I model is Border Patrol’s referral for prosecution under 8 U.S.C. 1325 or 1326, distinct from the U.S. Attorney’s subsequent filing decision and the court’s disposition, neither of which I observe. In practice the *referral* is the binding discretionary gate: prosecutors decline only about 1% of referred immigration cases, compared to roughly 17% of federal referrals overall (Motivans, 2022). To my knowledge, these records have not been used to study the internal structure of the referral decision. Rather than model prosecution with a single regression, I fit a sequence of gradient-boosted classifiers on rolling three-month windows and decompose each model with SHAP values, an approach from cooperative game theory

that attributes a prediction to its constituent features (Chen & Guestrin, 2016; Lundberg & Lee, 2017). This design lets me track, month by month, how much each observable characteristic of a case—age, prior apprehensions, criminal history, sector, region of origin—contributed to referral decisions as policy changed. It also captures nonlinearities and interactions that a coefficient-based decomposition would flatten.

This paper makes three contributions. Methodologically, it demonstrates a transferable design for detecting when and how a discretionary decision rule reorganizes, estimating that rule on a moving window and decomposing it without pre-specifying which features or interactions will matter (Bowman et al., 2024; Kleinberg et al., 2018). Substantively, it shows that Zero Tolerance did not switch discretion off and back on: referral rates peaked well below the universal mandate and largely receded, while the decision rule grew coarser and did not revert. Conceptually, it offers a lesson that travels beyond this case: in a policy environment where measures are enacted, enjoined, and rescinded in brief intervals, a temporary interruption can durably restructure the bureaucratic machinery it touches. That is, the formal status of a policy is a poor guide to its operational footprint.

Background

Criminal Prosecution at the Border

Criminal prosecution is one response available to the federal government when a person is apprehended crossing the border without authorization, but for much of modern immigration enforcement it was not the default (Eagly, 2020). Most apprehended migrants have been processed through the civil immigration system and removed administratively without a criminal charge (Kandel, 2021). Criminal prosecution is a categorically different sanction: it can involve federal imprisonment, produces a criminal record, and exposes the individual to enhanced penalties on any subsequent reentry (Eagly, 2010).

Criminal prosecution of immigration offenses dates to the Undesirable Aliens Act of 1929, which for the first time made unauthorized entry a federal crime rather than a purely civil matter (Lytle Hernández et al., 2021).¹ Two offenses, now codified at 8 U.S.C. 1325 and 1326, structure border prosecution. Section 1325 (improper entry) treats a first unauthorized entry as a misdemeanor punishable by up to six months; Section 1326 (reentry after removal) is a felony with penalties that escalate sharply when the defendant has a prior removal or criminal conviction (Eagly, 2020). Zero Tolerance directed prosecution specifically under Section 1325(a) (Eagly, 2020).

Although the statutes are old, their use as a routine, high-volume enforcement tool is recent. For most of the twentieth century, prosecution was rare relative to apprehensions; the federal government leaned on civil removal and reserved criminal charges for a small minority of cases. Beginning in the 2000s, however, immigration offenses grew to dominate the federal criminal

¹ The racially motivated Undesirable Aliens Act was passed as a compromise between Nativists concerned about the “Mexican problem” and industrialists who depended heavily on the Mexican migrant workforce (Lytle Hernández et al., 2021).

docket: improper entry and reentry together became the most frequently filed federal charges, and immigration cases came to account for an increasingly large share of all federal prosecutions (Light et al., 2014; Motivans, 2021).² Scholars of *crimmigration* describe this broader merger of criminal and immigration enforcement, in which criminal-justice tools are turned to the regulation of migration (Hernández, 2022); border prosecution is among its clearest expressions. This was the institutional setting in which Zero Tolerance arrived. By 2018, the federal system already had the courts, prosecutors, detention space, and legal template to prosecute border crossing at volume.

Operation Streamline and the Escalation of Prosecution

The most direct precursor to Zero Tolerance was Operation Streamline, a joint initiative of the Departments of Homeland Security and Justice launched in the Del Rio border sector in December 2005 (Corradini et al., 2018; Lydgate, 2010). Streamline was the first to adopt a self-described “zero-tolerance” approach to crossing in selected areas, referring apprehended migrants for criminal prosecution—often through consolidated, high-volume hearings—rather than processing them only through civil removal (Kimpel, 2022; Lydgate, 2010).³ Over the following years it spread sector by sector across the southwest border, and prosecution volumes rose accordingly, climbing from a few thousand cases annually in the early 2000s to the tens of thousands within a decade (Light et al., 2014; Lydgate, 2010).⁴

Two features of the Streamline era matter for understanding what came later. First, prosecution remained selective and uneven. The program was implemented sector by sector and never covered the entire border uniformly, so *where* a migrant was apprehended strongly conditioned whether prosecution was even on the table (Corradini et al., 2018). Second, agents and prosecutors retained discretion over whom to refer. As a matter of practice, first-time crossers, asylum seekers, and parents traveling with children were generally not sent to criminal court, precisely to avoid breaking up families (Isacson & Meyer, 2018; Lydgate, 2010). Streamline thus escalated prosecution without completely eliminating the judgment that selected cases for it.

Streamline was justified in part as a deterrent, yet evaluations found little evidence that criminal prosecution measurably reduced unauthorized crossing or re-crossing (Corradini et al., 2018). What is relevant for this analysis, however, is that by 2018 border referrals were still discretionary and uneven. Zero Tolerance is best understood as an attempt to override that discretion wholesale rather than as the creation of border prosecution from scratch.

² To illustrate the scope of this growth, in 1998, 63% of all federal arrests were of U.S. citizens. By 2018, this had inverted completely: 64% of all federal arrests were of non-U.S. citizens (Motivans, 2021, p. 1).

³ This was also the first time that a policy “significantly curtailed prosecutorial discretion—a hallmark of the U.S. criminal justice system—by mandating criminal prosecution of immigrants charged with entering or reentering the United States without authorization” (Corradini et al., 2018, p. 1)

⁴ In FY 2000 there were 19,500 federal immigration arrests in the five federal judicial districts along the U.S.–Mexico border. By FY 2008, that number had more than tripled to 70,583 (Motivans, 2021, p. 6).

Zero Tolerance: The Policy and Its Aftermath

The April 6, 2018 memorandum directed every U.S. Attorney's Office along the southwest border to adopt a zero-tolerance policy immediately for all offenses referred under Section 1325(a), superseding existing charging practices. The memo itself bound prosecutors, not Border Patrol; it applied only to cases DHS referred. Its implementation therefore depended on a parallel commitment from DHS to refer every amenable adult for prosecution, which the Department made in early May 2018 (DOJ OIG, 2021). In practice, then, the operative change fell on the referral decision this paper analyzes: Border Patrol was directed to stop selecting cases and refer them all. DHS determined that this policy would also cover adults traveling with minor children (DOJ OIG, 2021). Because that referral routed parents into criminal custody, where children cannot be held, the two were separated. Adults referred for prosecution were transferred to the U.S. Marshals Service; their children were reclassified as unaccompanied and placed with the Department of Health and Human Services, often without a reliable reunification mechanism. Later inspector-general reviews found that both departments moved forward despite knowing separations would follow and without adequate planning or interagency coordination (DHS OIG, 2018; DOJ OIG, 2021).

The political backlash came quickly. On June 20, 2018, President Trump issued an executive order that directed families to be kept together during proceedings, and within days a federal court ordered reunification for those already separated (Kandel, 2021). Neither action rescinded Zero Tolerance as a prosecution policy; both constrained family separation while leaving the prosecution mandate formally in place.⁵ However, although the memorandum was not formally withdrawn until January 2021, the zero-tolerance directive was effectively dead by late June 2018 (DHS OIG, 2018).⁶

The prosecution directive had a quieter afterlife than the policy itself, though a more ambiguous one than raw counts suggest. Entry-related prosecutions rose through fiscal year 2019, their peak, before pandemic-era measures reshaped enforcement and made later figures hard to compare (American Immigration Council, 2026; GAO, 2019). But counts track volume, not selection: because apprehensions were themselves rising, more prosecutions need not mean a higher per-apprehension referral rate, the quantity this paper measures. And the rate, as the analysis below shows, largely receded from its Zero Tolerance peak. The question is therefore not whether the rate stayed high—it did not—but whether the rule by which Border Patrol selected cases for referral reverted along with it. That is the empirical problem this paper takes up.

Prosecutorial Discretion and Its Measurement

Understanding that possibility requires situating border prosecution within the broader study of enforcement discretion. Discretion is a pervasive and consequential feature of criminal justice.

⁵ Recent work, however, establishes that the effects of family separation lasted long after its suspension (Amuedo-Dorantes & Bucheli, 2024).

⁶ The Department of Homeland Security's Office of Inspector General, reviewing the policy later in 2018, already described zero tolerance in the past tense and dated its operational end to the June 20 executive order: "...children were held in Border Patrol custody during the entire period the Zero Tolerance was in effect (May 5 to June 20, 2018)" (DHS OIG, 2018, p. 8).

Officials at every stage decide whom to charge, but those decisions are usually opaque—with little public scrutiny and only weakly constrained by formal rules (Bowman et al., 2024; Frederick & Stemen, 2012; Goldstein, 1963). Prosecutorial discretion is often understood as bounded decision-making under uncertainty. Prosecutors rely on observable case characteristics such as seriousness, the defendant’s record, and evidentiary strength as guides for charging decisions (Albonetti, 1987). More recent work has shown that both legal and extralegal defendant characteristics shape federal charging and declination decisions, and that their influence varies across districts and caseload conditions (Avni et al., 2024; Hartley & Tillyer, 2018; Shermer & Johnson, 2010). A policy that purports to *eliminate* discretion is, in this framing, asserting that those case characteristics should no longer matter.

Empirical studies usually measure discretion by decomposing decisions into the contributions of characteristics, often with regression models that estimate how each factor shifts the probability of an outcome (Bailey et al., 2025; Rehavi & Starr, 2014). That strategy is useful, but it can strain when the decision structure is nonlinear or when factors interact. For instance, prior work has found that case characteristics such as race, gender, and age condition one another, so that models omitting their interactions can understate their influence (Franklin, 2010). A single regression also averages the decision rule across an entire period. If Zero Tolerance changed which characteristics mattered and how they combined, a fixed functional form risks compressing that change into a handful of coefficient shifts. What is needed instead is a way to estimate the decision rule *over time*—to re-estimate it on a moving window so that its shape in one month can be compared directly to its shape a year later—and to decompose each estimate at the level of the individual case, so that interactions and nonlinearities are preserved rather than averaged away.

Recent work in the social sciences has shown that machine learning can illuminate exactly these decision structures, using flexible predictive models to analyze how human and institutional choices map case features onto outcomes (Kleinberg et al., 2018; Lovell et al., 2023). I adopt that strategy here, but with modest ambitions. My aim is not to advance a new theory of prosecutorial discretion, nor to adjudicate whether Zero Tolerance was sound policy. It is to provide a sharper empirical lens on a specific question: whether, and how, the logic by which border crossers were referred for prosecution was reorganized as enforcement moved from a baseline period through escalation, the Zero Tolerance directive, and its aftermath.⁷

Data and Methods

This study draws on individual-level apprehension data from the CBP FOIA Reading Room, a publicly accessible repository of records released by CBP in response to Freedom of Information Act requests. An apprehension is the “physical control or temporary detainment of a person who is not lawfully in the U.S. which may or may not result in an arrest” (CBP, 2024), and together, apprehensions are a reliable proxy for trends in undocumented immigration (Blanton, 2026;

⁷ For a fuller treatment of the visualization approaches available for SHAP-based model interpretation, see (Ponce-Bobadilla et al., 2024).

Espenshade, 1995). These apprehension data come from a series of “Encounter Subject Details” spreadsheets, each covering one fiscal year of Border Patrol “encounters” nationwide.⁸ Each file contains one record per apprehension event with information on demographics, apprehension context, and enforcement outcome. Because Zero Tolerance applied specifically to the U.S.–Mexico border, I restrict the data to apprehensions in the nine Southwest Border sectors.

The full extract covers fiscal years 2016 through the first quarter of fiscal year 2020—October 2015 through December 2019—and contains 1,960,075 apprehension records. I chose this window to center the Zero Tolerance period (April–June 2018) while retaining meaningful pre- and post-policy variation. It begins in the final year of the Obama administration and ends before Title 42 public health expulsions began in March 2020, a change that fundamentally altered the enforcement pipeline and removed most apprehended individuals from prosecution eligibility (Singer, 2022). Within this window, everyone in the sample was processed through the standard enforcement system and was, in principle, at risk of referral for prosecution.

From these records I draw an analytic sample of 1,000,781 Single Adult apprehensions, excluding Unaccompanied Children (219,251) and Family Unit Members (740,043). Unaccompanied children fall outside the 1325/1326 regime entirely: minors are not referred for criminal prosecution, and the Zero Tolerance directive did not apply to them. Family Unit Members require more care, because they are the group ostensibly at the center of the policy. Under Zero Tolerance, a parent traveling with a child who was referred for prosecution was separated from that child and processed as an adult defendant; in CBP records, such adults were reclassified from Family Unit Member to Single Adult, and their children reclassified as Unaccompanied (Amuedo-Dorantes & Bucheli, 2024; Kandel, 2021). The adults who actually experienced the Zero Tolerance prosecution regime are therefore already contained in the Single Adult sample. Those who remain coded as Family Unit Members are the parents who were *not* referred for prosecution. Among this group, 1325/1326 referral is almost never observed (a mean referral rate of 0.06% across all eras).

Outcome Variable

The primary outcome is whether an apprehended individual was referred for criminal prosecution, recorded in the FOIA data as a binary indicator (“Referred” / “Not referred”) under the field “Referred for Prosecution under 8 USC 1325 or 8 USC 1326.” Section 1325 criminalizes illegal entry as a misdemeanor (or a felony on subsequent offense); Section 1326 criminalizes illegal reentry after a prior removal as a felony.

Predictor Variables

The FOIA data include a set of individual-level characteristics that I used as predictor variables. Each was, in principle, observable to the Border Patrol agent at the time of apprehension and

⁸ In 2020, CBP updated the term to “encounters” which now includes not only apprehensions but also Title 42 expulsions—implemented during the COVID-19 pandemic—and Section 8 inadmissibles, referring to migrants denied entry due to legal grounds of inadmissibility (CBP 2024). In this paper I use the term apprehension because all the data I analyze was collected prior to 2020.

available to inform the referral decision. I organized these into four categories: demographic characteristics, enforcement history, apprehension context, and temporal indicators.

Demographic Characteristics

The data record age, sex, marital status, and citizenship country for each apprehended individual. I also grouped nationality into world regions, allowing the models to learn whether referral decisions respond to broad regional patterns, specific nationality signals, or both.

Enforcement History

Several variables capture prior contact with the enforcement system. The number of prior apprehensions records how many times the individual had previously been apprehended by Border Patrol. From the date of the earliest recorded apprehension I computed span of contact—days between that earliest encounter and the current one. I also computed recidivism velocity, defined as prior apprehensions per year of contact span. Together these distinguish a first-time crosser from a person with a long history of repeated apprehensions, a distinction central to prosecution decisions (Singer & Kandel, 2019). A Criminal Enforcement System (CES) indicator flags individuals with criminal history records in the National Crime Information Center (NCIC) database at apprehension and a drugs seized indicator records whether contraband was found during the encounter.

Apprehension Context

Each record includes the Border Patrol sector in which the apprehension occurred. The nine Southwest Border sectors differ in prosecution infrastructure, including federal court capacity and the presence of fast-track programs such as Operation Streamline. I also included time in the United States at apprehension, an ordinal variable that distinguishes recent crossers from individuals apprehended in the interior.

Temporal Indicators

I included calendar month and day of week. These temporal indicators absorb seasonal and within-period timing variation that would otherwise be attributed to correlated features such as sector and region of origin. I retain month, in particular, primarily to prevent the rolling-window models from loading transient temporal shifts onto substantively interpreted predictors.

Policy Eras

I organized the study period into four policy eras. The Obama baseline (January 2016–January 19, 2017) captures the final year of standard enforcement under the Obama administration. The Sessions escalation (January 20, 2017–April 5, 2018) begins with the Trump inauguration and covers the period during which Attorney General Sessions progressively directed increased immigration prosecution, before the formal Zero Tolerance mandate. The Zero Tolerance period (April 6–June 20, 2018) spans the memorandum through the executive order ending family separations. The post-

Zero Tolerance period (June 21, 2018–December 2019) follows the policy’s effective end. These eras structure the descriptive comparisons that follow, but they are not the unit of the main analysis: the rolling-window models treat the evolution of decision structures as continuous, allowing change to appear wherever it occurs rather than only at these pre-specified boundaries.

A Machine-Learning Approach to Decision Structure

The measurement problem laid out earlier—estimating a decision rule over time and decomposing it case by case—is poorly served by standard tools. A logistic regression predicting referral from case characteristics, perhaps stratified by era indicators, could offer interpretable coefficients and a familiar inferential framework. However, it assumes each predictor’s effect is linear and additive, requires interactions to be specified in advance, and imposes a single functional form across the whole period and population. For tracing how a multidimensional decision process may be reorganized across time, those assumptions are exactly what one needs to relax.

I therefore used gradient-boosted decision trees (XGBoost; Chen & Guestrin, 2016). The method builds an ensemble of shallow trees sequentially, each correcting the errors of the last. Because trees partition the data through if-then splits, the model captures nonlinear relationships, interactions, and subgroup-specific effects without the researcher specifying any of them; it discovers them from the data. The cost is interpretability: an ensemble of hundreds of trees yields no coefficient table. To recover interpretability I paired XGBoost with SHAP values (SHapley Additive exPlanations; Lundberg & Lee, 2017), which decomposes each individual prediction into additive contributions from each feature. Appendix B benchmarks this design against conventional logistic regression, fitting era-stratified and rolling-window logits on the same sample and predictors. The comparison shows that the main findings replicate in the familiar model and identifies what the flexible design adds: the discovery of the within-sector interaction collapse, which a regression recovers only once told which interactions to specify.

SHAP values derive from the “Shapley value” in cooperative game theory (Shapley, 1953), which fairly allocates a jointly produced payoff among contributors by averaging each one’s marginal contribution across *all* possible combinations of the others. In SHAP, the “players” are the predictors (age, criminal history, sector, nationality), the “game” is the model’s prediction for a given individual, and the “payoff” is how far that prediction departs from the *average* referral rate. For each person, SHAP assigns every predictor a value recording how much it pushed the predicted probability up or down relative to baseline, and these contributions sum exactly to the gap between that person’s prediction and the sample average. Averaging the *absolute* SHAP values across individuals then summarizes how strongly each feature shaped referral decisions overall: a high mean $|\text{SHAP}|$ marks a feature that moved predictions substantially, analogous to a large regression coefficient but without the assumption of linearity.

Rather than fit one model to the full period, I estimated the model on rolling windows and computed SHAP within each, producing a time series of decision structures. This allowed me to analyze not only which features mattered, but *when*—tracking how the referral rule’s reliance on each characteristic rose, fell, or reorganized as policy moved from baseline, through escalation, Zero Tolerance, and its aftermath.

Analytical Approach

The analysis proceeds from descriptive overview to progressively finer decompositions of the referral decision. I first report descriptive statistics for the analytic sample (Table 1) and plot monthly referral rates across the study period to document the aggregate policy shifts. Three model-based analyses then ask, in turn: how predictable referral was from observable characteristics and whether that predictability changed; which characteristics carried the most weight and how that shifted across eras; and how features interacted, and whether those interactions changed when policy did.

To trace the first two continuously, I trained XGBoost models on rolling three-month windows, sliding one month at a time across the study period. Each model trained on 80% of its window and was evaluated on the held-out 20%; from each I recorded the test AUC and the mean absolute value of SHAP ($|\text{SHAP}|$) for every predictor, yielding two time series—overall predictability and per-feature importance. AUC (the area under the ROC curve, ranging from 0.5 at chance to 1.0 at perfect discrimination) measures how well observable characteristics distinguish referred from non-referred individuals: a high value indicates structured, criteria-based selection; a low value indicates referral that is less predictable from observed features, whether because enforcement approached the blanket mandate (referring nearly everyone, so individual characteristics cease to discriminate) or because selection became more idiosyncratic. I fixed hyperparameters across all windows (300 trees, maximum depth 5, learning rate 0.05, 80% row and column subsampling).

Because the referral rate varies sharply across windows, I set *scale_pos_weight* in each window to its ratio of non-referred to referred cases, so that each model learned to discriminate referred cases even in windows where referral was rare. The rolling models show how importance evolved but do not yield individual-level values for closer inspection. I therefore fit a separate model to each of the four eras, using the same hyperparameters and predictors, and computed $|\text{SHAP}|$ values for a random subsample of up to 50,000 observations per era. This produced a person-level dataset of individuals with a $|\text{SHAP}|$ value for every feature, which supported dependence plots that show how a feature such as prior apprehensions or sector varies in effect across its range and across eras (akin to a marginal-effects plot, but free of the linearity a logistic model imposes), and direct comparison of feature rankings across eras.

Finally, because standard SHAP values treat each feature's contribution as independent, I computed SHAP interaction values (Lundberg et al., 2019) on each era model on up to 30,000 observations. For each feature pair the mean absolute interaction value summarizes how strongly the two jointly shape predictions beyond their separate effects, analogous to a regression interaction term but estimated for every pair at once and without a parametric form. Ranking pairs by interaction strength and comparing across eras reveals whether context-dependent associations were overridden when the mandate took hold.

Results

Prosecution Rates Across the Policy Eras

Table 1 reports descriptive statistics for the analytic sample of 1,000,781 Single Adult apprehensions recorded along the southern border between January 2016 and December 2019. The population is predominantly male (86%), with a mean age of 29, and originates overwhelmingly from Mexico (58%) and the Northern Central American countries of El Salvador, Guatemala, and Honduras (34%). Most apprehensions occur at entry (82%), and prior immigration histories are highly right-skewed (mean 1.18 prior apprehensions, SD 2.70). Criminal history (3%) and drug seizure (<1%) are uncommon. Apprehensions concentrate along the Texas border: the Rio Grande Valley sector alone accounts for 32% of the sample, with Tucson (17%), Laredo (12%), and San Diego (12%) next largest. Across the full period, 31% of these apprehensions were referred for federal criminal prosecution under 8 U.S.C. 1325 (illegal entry) or 1326 (illegal reentry).

[Table 1]

Because the analysis turns on change across enforcement periods, Table 2 disaggregates the sample by policy era. Its top row recaps the referral rate as an era average—20%, 29%, 45%, and 38% across the four periods. The rest of the table documents what otherwise changed about who was apprehended. The basic demographic profile is stable—age, sex, and marital status barely move—but several characteristics that bear on the referral decision shift markedly: the origin mix tilts away from Mexico (63% to 52%) toward Northern Central America and, late in the window, the Caribbean and South America; the population shifts toward first-time crossers (the zero-priors share rises from 53% to 60%); and recorded criminal history becomes scarcer (4.1% to 1.2%). These movements matter for interpretation: part of the structural change documented below could in principle reflect who was apprehended rather than how referral decisions were made, and Table 2 is the baseline against which that alternative is assessed.

Figure 1 plots the monthly prosecution referral rate over the four-year window, with policy periods marked. The aggregate 31% rate masks substantial temporal variation. The rate fell from 26% in January 2016 to approximately 17% by mid-year, where it remained for the rest of the Obama administration. After the change in administration referrals rose steadily—from 19% at the start of the Sessions escalation to a range of 27% to 33% through 2017 and into early 2018, standing at 26% in March 2018, the month before Zero Tolerance took effect. Under ZT the rate rose in two steps that track the directive's implementation: 29% in April 2018, when the DOJ memo was in place but DHS had not yet formally committed to referring all adults, then 52% in May, the first month of the full referral mandate, and 59% in June (the final ZT month). It held at 59% in July and reached its peak of 61% in August 2018—two months after the executive order ending the policy—before declining steadily and returning to the Sessions-escalation baseline (~30%) by March 2019, holding there through the rest of the observation window.

[Figure 1]

There are two important patterns to note here. First, even at its peak, the referral rate never approached the 100% target: a maximum monthly rate of 61% indicates that prosecutorial discretion was substantially compressed but not eliminated, even under the most aggressive enforcement period in the data. Second, the elevation persisted well beyond the three-month ZT window. The referral rate did not decline to pre-ZT levels until about seven months after the August peak.

Prior apprehension history has long been among the strongest predictors of prosecution referral—partly mechanical, since 8 U.S.C. 1326 reentry charges require a documented prior immigration encounter, and partly discretionary, with prosecutors and referring agents historically prioritizing repeat crossers (Abrego et al., 2017; Argueta, 2016). As such, Figure 2 disaggregates monthly prosecution rates by prior apprehension count across the four policy periods.

[Figure 2]

Pre-ZT, prosecution rates rose monotonically with prior apprehension history: 21.8% for first-time apprehensions, 33.2% for one prior, 38.3% for two or three priors, and 42.3% for four or more. Under Zero Tolerance rates rose across all groups, but the rise was inversely related to baseline. First-time apprehensions saw a 2.4-fold increase (to 52.9%); the higher-priors groups saw progressively smaller increases—1.8-fold for one prior (to 59.8%), 1.5-fold for two or three priors (to 58.8%), and 1.4-fold for four or more (to 58.2%). At peak ZT the four groups compressed into a narrow band of roughly 53–60%: first-time crossers remained distinctly lower than the rest, but the three priors-bearing groups converged to nearly identical rates, shrinking the gap between the lowest- and highest-priors groups from 20.5 to 5.3 percentage points. The directive thus substantially flattened—though did not eliminate—prior-history-based selectivity.

These elevated rates did not persist, though they faded slowly. Rates remained near their Zero Tolerance peak through the second half of 2018—first-time apprehensions still averaged 48.3% from July through December—then fell steadily over 2019. By the second half of that year, every priors group sat within two points of its January–March 2018 level, and the gap between the lowest- and highest-priors groups, compressed to 5.3 points at peak Zero Tolerance, had reopened to 20.4 points, restoring the pre-ZT gradient at its full width. The priors-conditioned rate structure thus reverted to its Sessions-era baseline within roughly a year of the directive's end. What did not revert was the underlying decision rule, whose reorganization the rolling-window analysis documents below.

How the Decision Structure Changed

To trace how the structure of prosecution decisions evolved I trained XGBoost models on rolling three-month windows sliding across the study period one month at a time, recording the area under the ROC curve (AUC) on a 20% holdout for each window as a measure of how predictable prosecution is given the full feature set. A high AUC indicates that observable features collectively explain who is referred; a lower AUC indicates that prosecution decisions are less predictable from those features. Figure 3 plots the rolling AUC across the study period.

[Figure 3]

Pre-ZT rolling AUCs sit between 0.88 and 0.93 with no meaningful drift, indicating prosecution referral was highly structured and readily predictable from observed case features. The series begins to decline as the rolling window first overlaps with ZT months, reaching 0.86 in the window that fully spans April through June 2018, dropping to 0.80 in the subsequent window, and then stabilizing between 0.80 and 0.82 through the rest of the observation period. The 0.10-point drop is modest in absolute terms but substantively larger as a change in discriminatory capacity: pre-ZT capacity above chance averaged 0.40 ($= 0.90 - 0.50$), post-ZT 0.31 ($= 0.81 - 0.50$), a reduction of approximately 22%. A meaningful portion of the order that previously governed prosecution referral was lost during ZT and *not* recovered a year and a half later. The decision structure that governed referral pre-ZT did not survive the policy episode: even after rates returned to the Sessions-escalation baseline, the underlying selection function remained measurably less ordered. The AUC decline tracks ZT's entry into the rolling window rather than preceding it, which attributes the structural disruption to ZT itself rather than to a broader pre-existing trend. This pattern indicates the post-ZT enforcement environment operates under a reorganized prosecution decision rule.

The AUC trajectory shows that the prosecution decision function became less predictable but does not show *which* features drove the change. To examine how the model's reliance on individual features evolved, I extracted the mean absolute SHAP value for each predictor in each rolling window. Figure 4 plots these trajectories with sector—the dominant feature throughout—separated into its own panel. Figures 4 and 5 display the ten features with the largest mean $|\text{SHAP}|$ averaged across all windows; the remaining predictors are aggregated as “Other” in Figure 5. Two interpretive points warrant emphasis. First, mean $|\text{SHAP}|$ reflects both a feature's predictive contribution and the model's overall predictive capacity in a given window. For instance, as AUC declined post-ZT, the total predictive signal across all features contracted by roughly 37% (pre-ZT cross-window mean 3.35, post-ZT 2.10). A generalized downward shift in absolute $|\text{SHAP}|$ is therefore expected and does not by itself indicate that any one feature became less informative relative to the others. Interpretation accordingly draws on two complementary readings: absolute trajectories, which capture how the model's overall reliance on each feature changed, and shares of total $|\text{SHAP}|$, which capture how predictive weight was redistributed across features as the total contracted.

[Figure 4]

Sector dominates the predictive signal throughout, but its trajectory mirrors the AUC story. Pre-ZT, sector's mean $|\text{SHAP}|$ ranges between roughly 1.02 and 1.55, peaking in late 2017 during the Sessions escalation. As ZT enters the rolling window, sector's contribution falls sharply to 0.88 in the window fully spanning ZT and to 0.69 in the subsequent window. Post-ZT it continues a slow decline, ending the observation period at 0.60. The cross-window mean falls from 1.27 pre-ZT to 0.71 post-ZT, a 44% drop. Sector's share of total signal declines more modestly, from 38.0% to 34.1%, indicating its dominance eroded somewhat even as it remained the most informative feature by a wide margin.

The bottom panel reveals two distinct patterns among the non-sector features. Features capturing characteristics of the individual contracted sharply in both absolute and relative terms. Age fell from 0.41 pre-ZT to 0.12 post-ZT (−71%); its share of total signal more than halved from 12.3% to 5.7%, the largest share loss of any feature. Prior apprehensions fell from 0.17 to 0.03 (−82%; 5.0% → 1.4%); criminal history collapsed from 0.13 to 0.02 (−86%; 3.8% → 0.9%). Span of contact—time elapsed between the individual’s earliest Border Patrol apprehension and the current one—fell from 0.28 to 0.12 in absolute terms and from 8.4% to 5.5% in share. Among the priors-related features, only recidivism velocity (apprehensions per year of contact) retained its proportional weight, holding a stable share of roughly 7–8% across periods even as its absolute value declined.

Two contextual features moved in the opposite direction. World region of origin rose from 0.15 pre-ZT to 0.23 post-ZT (+54%), more than doubling its share of total signal from 4.4% to 10.9%—the largest share gain of any feature, and the only feature in the model whose absolute |SHAP| meaningfully increased post-ZT. Day-of-week also gained share (5.3% → 9.4%) despite a near-stable absolute trajectory; its share rose not because day of week grew in influence but because the case-specific features lost absolute signal, so its roughly constant contribution came to occupy a larger fraction of a smaller total. Sex (female) shows a similar pattern (3.9% → 6.2%) without absolute increase.

Month-of-year exhibits a third pattern. Rather than rising gradually, month’s mean |SHAP| spikes sharply in two specific windows—those containing the April-June 2018 prosecution-rate jump, and those containing the early 2019 decline back toward baseline—reaching peaks of approximately 0.68 before returning to near-baseline. These spikes coincide with periods in which prosecution rates were changing rapidly within the three-month window, making month a mechanically strong proxy for early-vs-late portion of the window. Outside the transition windows, month’s contribution is comparable to its pre-ZT level. I therefore read the month spikes as a diagnostic marker of regime instability rather than a stable post-ZT increase: month becomes informative precisely at the moments when the prosecution decision rule was shifting most rapidly, which is consistent with the AUC trajectory’s analogous breakpoints at the same moments.

Figure 5 displays the same data as relative shares stacked to 100% in each window. The reorganization is more directly visible: the bottom of the stack—sector and the case-specific features—contracts visibly after April 2018, while the top—world region especially—expands. The two month-of-year share spikes appear as transient pink incursions during the regime-transition windows.

[Figure 5]

The mean |SHAP| values in Figures 4 and 5 implicitly combine two sources of importance: a feature’s main effect, or its independent contribution, and its interaction effects, or the portion of its contribution that depends on other features. To examine the interaction component directly, I computed SHAP interaction values, which decompose each feature’s overall contribution into a main-effect component and pairwise components with every other feature. Pairwise computation is substantially more computationally intensive than standard SHAP, so I estimated these values on

era-specific models trained on the full set of observations within each era rather than on rolling windows. Figure 6 displays the mean absolute interaction value for the ten feature pairs with the largest cross-era magnitudes.

[Figure 6]

The dominant pattern in Figure 6 is the collapse of sector $\times$ case-feature interactions during and after ZT. Age $\times$ Sector falls from 0.095 in the Obama era and 0.092 during Sessions to 0.014 both during ZT and post-ZT, which represents an 85% decline in the joint contribution of age and sector to prosecution decisions. Span of contact $\times$ Sector follows a similar trajectory (0.078 $\rightarrow$ 0.039 $\rightarrow$ 0.018 $\rightarrow$ 0.013), as does Prior apprehensions $\times$ Sector (0.047 $\rightarrow$ 0.055 $\rightarrow$ 0.009 $\rightarrow$ 0.002), with the latter effectively eliminating its joint contribution by post-ZT. Recidivism velocity $\times$ Sector declines through ZT with a partial recovery post-ZT (0.069 $\rightarrow$ 0.053 $\rightarrow$ 0.025 $\rightarrow$ 0.031). Two non-sector interactions among case features collapse nearly completely: Age $\times$ Recidivism velocity from 0.056 to 0.006 (-90%) and Age $\times$ Span of contact from 0.038 to 0.003 (-91%). Three interactions move in the opposite direction. Sector $\times$ Day of week rises sharply during Sessions and ZT (0.037 $\rightarrow$ 0.105 $\rightarrow$ 0.104) before partially declining post-ZT (0.072); Sector $\times$ Time in US peaks during ZT (0.057); and World region $\times$ Sector follows a U-shaped trajectory (0.044 $\rightarrow$ 0.033 $\rightarrow$ 0.023 $\rightarrow$ 0.033), partially recovering post-ZT.

This pattern provides a more specific account of the SHAP shifts documented above. Sector’s overall mean $|\text{SHAP}|$ dropped substantially in absolute terms (1.27 $\rightarrow$ 0.71, a 44% decline) while its share of total signal fell only modestly (38.0% $\rightarrow$ 34.1%). The interaction analysis indicates that sector’s role changed in kind, not merely in magnitude: pre-ZT, much of sector’s predictive contribution came from joint contributions with case-level features, with sectors weighing age, prior history, span of contact, and recidivism velocity differently in arriving at prosecution decisions. Post-ZT, those joint contributions largely collapsed, leaving sector’s contribution closer to a pure main effect—which sector an apprehension occurred in still matters substantially, but how case characteristics modulate prosecution within sectors has flattened.

The post-ZT enforcement environment thus operated with reduced reliance on case-level features overall and substantially less integration between case-level features and the sectoral context in which the prosecution decision is made. Appendix Figures A1 and A2 visualize this sector-level flattening directly, showing how sectors’ previously divergent age and prior-apprehension responses converged toward uniform near-zero patterns by ZT and remained flat thereafter.

Taken together, the AUC, rolling SHAP, and interaction analyses describe a coherent structural transformation of the prosecution decision function. Pre-ZT, the decision was substantially individuated within sectors: location dominated, but case-specific features (age, prior history, criminal background, span of contact) collectively carried a large portion of the predictive signal. Moreover, sector itself derived much of its importance from joint contributions with these case-level features. In the post-ZT period, this individuated structure largely collapsed. Sector remained dominant in share but functioned more as a flat categorical sort than as a context that modulates the weight of case features. Case-specific features lost both absolute and relative importance; world region of origin grew substantially in importance; and the case-by-case calibration

that previously occurred within sectors flattened. The post-ZT decision function depended increasingly on broad attributes—where one crossed and where one came from—and decreasingly on the case-specific features and within-sector calibrations that previously informed referral decisions. This pattern is consistent with the AUC contraction: categorical features are inherently coarser predictors than the individuated case factors and within-sector calibrations they displaced, and an environment that relies more on categorical sorting will, almost by construction, generate decisions that are less predictable from observed case features.

Discussion

This paper asked whether Zero Tolerance eliminated prosecutorial discretion at the border, and what became of that discretion after the policy ended. Using 1,000,781 Single Adult apprehensions from January 2016 through December 2019, and a sequence of rolling-window models decomposed with SHAP values, it traced the referral decision before, during, and after the directive, measuring not only how often Border Patrol referred cases but the rule by which it selected them. A universal referral mandate is, by definition, the elimination of prosecutorial discretion: if every apprehended adult is referred, no case-specific judgment remains. Yet under Zero Tolerance discretion was never eliminated. Even at its August 2018 peak, the monthly referral rate reached only 61%, so individual characteristics still sorted cases in the most aggressive enforcement month on record. Discretion survived, compressed, but what changed was its internal logic.

The post-ZT regime did not restore the pre-ZT selection logic. It returned a reorganized decision function in which where a person crossed and where they came from did most of the sorting, while the case-specific features that had previously calibrated referral within sectors, namely age, prior history, criminal record, and span of contact, lost substantial weight. World region of origin was the only feature whose absolute contribution to referral rose after the policy, a 54% increase, with its share of total predictive signal almost tripling from 4.4 to 11%, the largest gain of any feature. The case-specific features moved the other way. Criminal history fell 86% in absolute terms, from 3.8 to 0.9% of total signal, and age, prior apprehensions, and span of contact contracted alongside it. Each pattern amplifies the other's normative weight, because a referral rule that relies more on where a person was born, a broad categorical attribute, and less on criminal history, a legally relevant factor the statutes themselves single out, is troubling from both directions at once.⁹

The most consequential feature of this reorganization is that it did not end when the policy did. Zero Tolerance was effectively dead by late June 2018 and not formally withdrawn until January 2021; roughly eighteen months after its operational end, the referral structure it induced was largely still in place. The policy's legal timeline and its operational timeline came apart. What the memo built—surge prosecutorial capacity, default referral heuristics, case-management routines—did not

⁹ One qualification matters here. Recorded criminal history also grew rarer in the apprehended population over the study period, from 4.1 to 1.2% (Table 2). A logistic-regression benchmark (Appendix B) helps separate the two effects: criminal history's marginal effect on an individual who has a record fell only modestly, while its aggregate contribution to the model collapsed largely because the feature itself became scarce. The decline in its predictive footprint is thus as much a change in who was apprehended as a change in how referral decisions weighted a criminal record.

recede when the memo's force did, and none of it required a standing directive to persist. Eleven weeks of mandate were enough to reset the baseline, and there was no memo to rescind that would have reset it back.

Recovering this required a method that could distinguish a change in the rate of referral from a change in the rule generating it, a distinction the standard toolkit tends to average away. Rolling-window XGBoost decomposed with SHAP values estimates the referral rule locally in time and attributes each decision case by case. A conventional model recovers much of this: as a logistic-regression benchmark shows (Appendix B), a logit fit on the same rolling windows reproduces both the decline in predictability and the broad reweighting toward region and away from case-specific features. What it does not deliver on its own is the reorganization's most distinctive feature—the collapse of within-sector calibration, which surfaces only once one already knows to specify the relevant interactions—and a single importance scale comparable across features of different types. The tree ensemble searches the space of interactions without being told, and SHAP places every feature on one additive scale, so the collapse emerges as a finding rather than a hypothesis one had to anticipate.

Three limitations bound the findings. First, the structural series is right-censored at March 2020, when Title 42 public health expulsions removed most apprehended individuals from prosecution eligibility and closed the pipeline measured here, so whether the post-ZT structure would have persisted, drifted, or eventually reverted is unobservable. Second, the data capture the Border Patrol referral gate rather than DOJ filing or disposition. While the roughly 1% declination rate makes referral the binding gate, downstream filtering could differ in ways these records cannot show. And SHAP describes the rule only as expressed through recorded characteristics, so selection that shifted partly onto unobserved criteria would surface as reduced predictability from observables. Some of the AUC contraction may reflect this. It cannot, however, account for the full pattern: an unobserved shift would dim the observable signal, not reorganize it, and the collapse of specific sector-by-case-feature interactions and the reweighting toward region are changes in the structure of what remains observed.

Zero Tolerance ran eleven weeks as an operational directive, and the decision structure it left behind was still in place a year and a half later. Aggregate referral rates did not register this. They receded toward their earlier range while the rule beneath them stayed reorganized, which is why the change is visible only from inside the decision. Discretion at the border was not eliminated and did not return as it had been. Through the end of the observation window, it persisted in a coarser form. Taken together, these findings suggest that the policies that matter most for long-run institutional behavior may be the ones that look briefest on paper

Data Availability Statement

The data and code required to reproduce all results are openly available in an anonymized repository for peer review at https://osf.io/hgytk/overview?view_only=123dceb09c6b485f9ee4fd0f1d477267. The analysis uses individual-level apprehension records released by U.S. Customs and Border Protection (CBP) under the Freedom of Information Act and obtained from the CBP Documents Library (<https://www.cbp.gov/document/foia-record/customs-and-border-protection-border->

[patrol-statistics](#)). As works of the U.S. federal government, these records are in the public domain. The repository provides the combined source data, all data-preparation and analysis code (Stata and Python), and instructions to reproduce every table and figure. A permanent, citable archive with a DOI will replace the anonymized link upon acceptance.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, I used Claude Opus 4.8 (Anthropic) for copyediting. After using this tool, I reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

- Abrego, L., Coleman, M., Martínez, D. E., Menjívar, C., & Slack, J. (2017). Making Immigrants into Criminals: Legal Processes of Criminalization in the Post-IIRIRA Era. *Journal on Migration and Human Security*, 5(3), 694–715. <https://doi.org/10.1177/233150241700500308>
- Albonetti, C. A. (1987). Prosecutorial Discretion: The Effects of Uncertainty. *Law & Society Review*, 21(2), 291–313. <https://doi.org/10.2307/3053523>
- American Immigration Council. (2026). *Prosecuting People for Coming to the United States* [Fact Sheet]. American Immigration Council. <https://www.americanimmigrationcouncil.org/fact-sheet/immigration-prosecutions/>
- Amuedo-Dorantes, C., & Bucheli, J. R. (2024). Family Separation and Reunification Under President Trump's Zero-Tolerance Policy. *International Migration Review*, 58(2), 989–1016. <https://doi.org/10.1177/01979183231162619>
- Argueta, C. N. (2016). *Border Security: Immigration Enforcement Between Ports of Entry* (Nos. 7–5700). Congressional Research Service. https://www.fosterglobal.com/policy_papers/2016/bordersec.pdf
- Avni, O., Guetzkow, J., & Hasisi, B. (2024). Bias in prosecutorial decision making: Bridging focal concerns & group threat. *Journal of Criminal Justice*, 92, 102192. <https://doi.org/10.1016/j.jcrimjus.2024.102192>
- Bailey, C. M., Hylton, S., Henry, T. K. S., Diaz, C. L., & Grommon, E. (2025). Exploring racial disparities across prosecutorial decision points: Evidence from a Midwest prosecutor's office. *Journal of Crime and Justice*, 0(0), 1–19. <https://doi.org/10.1080/0735648X.2025.2483329>

- Blanton, M. (2026). Brief Report: Are Customs and Border Protection (CBP) Apprehensions a Reliable Proxy for Undocumented Migration? New Evidence on Migration from Guatemala. *Population Research and Policy Review*, 45(3), 29. <https://doi.org/10.1007/s11113-026-10012-8>
- Bowman, R., Lowrey-Kinberg, B., & Gould, J. (2024). An integrated model of prosecutor decision-making. *Law & Society Review*, 58(3), 452–480. <https://doi.org/10.1017/lsr.2024.24>
- CBP. (2024). *CBP Enforcement Statistics*. <https://www.cbp.gov/newsroom/stats/cbp-enforcement-statistics>
- Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785>
- Corradini, M., Kringen, J. A., Simich, L., Berberich, K., & Emigh, M. (2018). *Operation Streamline: No Evidence that Criminal Prosecution Deters Migration*. Vera Institute of Justice.
- DHS OIG. (2018). *OIG-18-84—Special Review—Initial Observations Regarding Family Separation Issues Under the Zero Tolerance Policy* (OIG-18-84). Department of Homeland Security Office of the Inspector General. <https://www.oig.dhs.gov/sites/default/files/assets/2018-10/OIG-18-84-Sep18.pdf>
- DOJ OIG. (2021). *Review of the Department of Justice’s Planning and Implementation of Its Zero Tolerance Policy and Its Coordination with the Departments of Homeland Security and Health and Human Services* (Nos. 21–028). Department of Justice Office of the Inspector General. <https://oig.justice.gov/reports/review-department-justices-planning-and-implementation-its-zero-tolerance-policy-and-its>
- Eagly, I. (2010). Prosecuting Immigration. *Northwestern University Law Review*, 104(4), 1281–1360. <https://papers.ssrn.com/abstract=1710182>
- Eagly, I. (2020). The Movement to Decriminalize Border Crossing. *Boston College Law Review*, 61, 1967–2029.
- Espenshade, T. J. (1995). Using INS Border Apprehension Data to Measure the Flow of Undocumented Migrants Crossing the U.S.-Mexico Frontier. *International Migration Review*, 29(2), 545–565. <https://doi.org/10.1177/019791839502900209>
- Franklin, T. W. (2010). The intersection of defendants’ race, gender, and age in prosecutorial decision making. *Journal of Criminal Justice*, 38(2), 185–192. <https://doi.org/10.1016/j.jcrimjus.2009.12.001>
- Frederick, B., & Stemen, D. (2012). *The Anatomy of Discretion: An Analysis of Prosecutorial Decision Making* [Technical Report]. Vera Institute of Justice.
- GAO. (2019). *Immigration Related Prosecutions Increased from 2017 to 2018 in Response to U.S. Attorney General’s Direction* (GAO-20-172). United States Government Accountability Office. <https://www.gao.gov/products/gao-20-172>

- Goldstein, H. (1963). Police Discretion: The Ideal versus the Real. *Public Administration Review*, 23(3), 140–148. <https://doi.org/10.2307/973838>
- Hartley, R. D., & Tillyer, R. (2018). Examining Prosecutorial Discretion in Federal Criminal Cases: Legal and Extra-Legal Determinants of Declination and Charge Change Decisions. *Justice Quarterly*, 35(7), 1195–1225. <https://doi.org/10.1080/07418825.2018.1530367>
- Hernández, C. C. G. (2022). *Crimmigration Law* (Second). American Bar Association.
- Isacson, A., & Meyer, M. (2018). *WOLA Report: The Zero Tolerance Policy*. Washington Office on Latin America. <https://www.wola.org/analysis/wola-report-zero-tolerance-policy/>
- Kandel, W. A. (2021). *The Trump Administration's "Zero Tolerance" Immigration Enforcement Policy* (No. R45266). Congressional Research Service.
- Kimpel, A. F. (2022). Alienating Criminal Procedure. *Georgetown Immigration Law Journal*, 37, 237–289.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., & Mullainathan, S. (2018). Human Decisions and Machine Predictions. *The Quarterly Journal of Economics*, 133(1), 237–293. <https://doi.org/10.1093/qje/qjx032>
- Light, M. T., Lopez, M. H., & Gonzalez-Barrera, A. (2014, March 18). The Rise of Federal Immigration Crimes. *Pew Research Center*. <https://www.pewresearch.org/race-and-ethnicity/2014/03/18/the-rise-of-federal-immigration-crimes/>
- Lovell, R. E., Klingenstein, J., Du, J., Overman, L., Sabo, D., Ye, X., & Flannery, D. J. (2023). Using machine learning to assess rape reports: “Signaling” words about victims’ credibility that predict investigative and prosecutorial outcomes. *Journal of Criminal Justice*, 88, 102107. <https://doi.org/10.1016/j.jcrimjus.2023.102107>
- Lundberg, S. M., Erion, G. G., & Lee, S.-I. (2019). *Consistent Individualized Feature Attribution for Tree Ensembles* (arXiv:1802.03888). arXiv. <https://doi.org/10.48550/arXiv.1802.03888>
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, 4768–4777. <https://dl.acm.org/doi/10.5555/3295222.3295230>
- Lydgate, J. J. (2010). Assembly-Line Justice: A Review of Operation Streamline. *California Law Review*, 98(2), 481–544. <https://www.jstor.org/stable/20743978>
- Lytle Hernández, K., Ngai, M., & Eagly, I. (2021). *Brief for Professors Kelly Lytle Hernández, Mae Ngai, and Ingrid Eagly as Amici Curiae Supporting Respondent, United States v. Refugio Palomar-Santiago*. U.S. Supreme Court.
- Motivans, M. (2021). *Immigration, Citizenship, and the Federal Justice System, 1998–2018* (NCJ 253116). U.S. Department of Justice Bureau of Justice Statistics. <https://bjs.ojp.gov/content/pub/pdf/icfs9818.pdf>

- Motivans, M. (2022). *Federal justice statistics, 2020* (NCJ 304314). U.S. Department of Justice Bureau of Justice Statistics. <https://doi.org/10.1037/e542272008-001>
- Ponce-Bobadilla, A. V., Schmitt, V., Maier, C. S., Mensing, S., & Stodtmann, S. (2024). Practical guide to SHAP analysis: Explaining supervised machine learning model predictions in drug development. *Clinical and Translational Science*, 17(11), e70056. <https://doi.org/10.1111/cts.70056>
- Rehavi, M. M., & Starr, S. B. (2014). Racial Disparity in Federal Criminal Sentences. *Journal of Political Economy*, 122(6), 1320–1354. <https://doi.org/10.1086/677255>
- Shapley, L. S. (1953). A value for n-person games. In H. W. Kuhn (Ed.), *Contributions to the Theory of Games*. Princeton University Press.
- Shermer, L. O., & Johnson, B. D. (2010). Criminal Prosecutions: Examining Prosecutorial Discretion and Charge Reductions in U.S. Federal District Courts. *Justice Quarterly*, 27(3), 394–430. <https://doi.org/10.1080/07418820902856972>
- Singer, A. (2022). *U.S. Border Patrol Apprehensions and Title 42 Expulsions at the Southwest Border: Fact Sheet*. Congressional Research Service. <https://crsreports.congress.gov/product/pdf/R/R47343>
- Singer, A., & Kandel, W. A. (2019). *Immigration: Recent Apprehension Trends at the U.S. Southwest Border*. Congressional Research Service. https://tracfed.syr.edu/tracker/dynadata/2019_12/R46012.pdf
- TRAC Immigration. (2018). “Zero Tolerance” at the Border: Rhetoric vs. Reality (No. 520). Transactional Records Access Clearinghouse. <https://tracreports.org/immigration/reports/520/>
- U.S. Department of Justice. (2018, April 6). *Attorney General Announces Zero-Tolerance Policy for Criminal Illegal Entry*. The United States Department of Justice Office of Public Affairs. <https://www.justice.gov/archives/opa/pr/attorney-general-announces-zero-tolerance-policy-criminal-illegal-entry>

Tables and Figures

Table 1. Descriptive Statistics, Single Adult Apprehensions, Jan 2016 – Dec 2019

Variable	Mean or %	SD
Referred for prosecution (1325/1326)	31.1	
Age at apprehension	29.40	9.08
Number of previous apprehensions	1.18	2.70
Days from earliest apprehension to current	2482.74	2326.20
Prior apprehensions per year of contact	4.07	6.92
Sex (female)	14.1	
Criminal history (CES)	2.6	
Drugs seized	0.3	
World region		
Mexico	58.2	
Northern Central America	33.5	
Other Central America	1.0	
South America	2.3	
Caribbean	1.3	
Asia	3.4	
Middle East & Africa	0.2	
Europe	0.1	
Other/Unknown	0.0	
Apprehension sector		
Big Bend (BBT)	2.1	
Del Rio (DRI)	6.0	
El Centro (ELC)	8.1	
El Paso (EPT)	8.0	
Laredo (LRT)	12.0	
Rio Grande Valley (RGV)	32.3	
San Diego (SDC)	12.3	
Tucson (TCA)	16.8	
Yuma (YUM)	2.4	
Marital status		
Married	12.1	
Single	83.5	
Divorced	0.2	
Separated	0.1	
Widowed	0.1	
Unknown	4.0	
Time in US at apprehension		
At Entry	82.1	
Within 72 hours	7.4	
4 to 30 days	7.9	
1 month to 1 year	0.6	
Over 1 year	1.6	
Not Applicable	0.4	
N	1,000,781	


Note: Single Adult apprehensions only (N = 1,000,781). Continuous variables are summarized by mean and SD; all other entries are percentages of the indicated category. Percentages may not sum to 100 due to rounding.

Table 2. Sample Composition by Policy Era, Single Adult Apprehensions, Jan 2016–Dec 2019

Variable	Obama baseline	Sessions escalation	Zero Tolerance	Post-ZT
Referred for prosecution (%)	19.9	29.1	45.3	37.7
World region (%)				
Mexico	63.1	64.5	54.9	52.2
Northern Central America	31.6	30.0	37.9	35.9
Other Central America	0.4	0.3	0.6	1.7
South America	1.8	1.3	1.5	3.3
Caribbean	0.2	0.2	0.1	2.8
Asia	2.6	3.4	4.7	3.8
Middle East & Africa	0.1	0.1	0.2	0.2
Europe	0.1	0.1	0.1	0.1
Other/Unknown	0.0	0.1	0.0	0.0
Prior apprehensions (%)				
0 priors	53.2	53.1	57.9	59.6
1 prior	20.7	21.1	20.1	19.7
2–3 priors	15.9	15.7	14.0	13.6
4+ priors	10.3	10.1	8.1	7.2
Criminal history, CES (% yes)	4.1	3.8	1.8	1.2
Apprehension sector (%)				
Big Bend (BBT)	1.6	3.0	2.3	2.0
Del Rio (DRT)	6.0	5.0	4.7	6.8
El Centro (ELC)	6.5	9.4	9.3	8.3
El Paso (EPT)	6.4	6.1	6.2	10.3
Laredo (LRT)	11.4	12.5	11.9	12.0
Rio Grande Valley (RGV)	36.6	30.3	33.1	30.4
San Diego (SDC)	10.5	13.4	12.3	13.0
Tucson (TCA)	19.3	18.1	16.7	14.5
Yuma (YUM)	1.8	2.2	3.5	2.7
N	288,577	221,598	59,079	431,527


Note: Single Adult apprehensions only (N = 1,000,781). Cell entries are within-era (column) percentages. Obama baseline (Jan 2016–Jan 19, 2017), Sessions escalation (Jan 20, 2017–Apr 5, 2018), Zero Tolerance (Apr 6–Jun 20, 2018), Post-ZT (Jun 21, 2018–Dec 2019). Columns may not sum to 100 due to rounding.

Figure 1. Monthly Prosecution Referral Rate, Jan 2016 – Dec 2019




Note: Monthly share of Single Adult apprehensions referred for prosecution under 8 U.S.C. 1325 or 1326. Shaded bands mark the four policy eras.

Figure 2. Monthly Prosecution Rate by Prior Apprehension History




Note: Monthly referral rate computed within each prior-apprehension group (0, 1, 2–3, and 4+ priors), Single Adult apprehensions. Shaded bands mark the four policy eras.

Figure 3. Rolling Three-Month AUC, Jan 2016 – Dec 2019




Note: XGBoost models trained on rolling three-month windows of Single Adult apprehensions, sliding one month at a time. For each window, the model is trained on 80% of observations and AUC is computed on the 20% holdout. Dashed gray reference line at 0.5 (random guessing).

Figure 4. Rolling Three-Month Mean $|\text{SHAP}|$ by Feature, Jan 2016 – Dec 2019




Note: Two-panel plot. Top panel: sector. Bottom panel: nine other model features (age, recidivism velocity, span of contact, day of week, world region, month of year, female, prior apprehensions, criminal history); shown are the ten features with the largest mean $|\text{SHAP}|$ averaged across all windows. Same rolling three-month windows as Figure 3.

Figure 5. Share of Total Predictive Signal by Feature, Jan 2016 – Dec 2019



Note: Stacked area chart, 100% normalized in each window. Same rolling three-month windows as Figures 3 and 4. The named features are the ten with the largest mean $|SHAP|$ across all windows; remaining predictors are aggregated as Other. Each feature's color matches its color in Figure 4. "Other" combines small stable features (nationality, marital status, drugs seized, time in US). Stacking order (bottom to top): sector, age, span of contact, recidivism velocity, prior apprehensions, criminal history, female, Other, day of week, month of year, world region.

Figure 6. SHAP Interaction Strength Across Policy Eras



Note: Heatmap of mean absolute SHAP interaction values for the ten feature pairs with the largest cross-era magnitudes. Values computed from era-specific XGBoost models trained on Single Adult apprehensions in each policy era. Cells are shaded by magnitude; columns are policy eras (Obama, Sessions, ZT, Post-ZT).