Anticipating State Action: Risk Perceptions and Consumption under Immigration Enforcement*

Alberto Ciancio[†1] and Camilo García-Jimeno[‡2]

[1]University of Glasgow

[2]Federal Reserve Bank of Chicago

April 21, 2026

Abstract

Immigration enforcement affects millions of individuals who make daily economic decisions under uncertainty about state action. Using data from 2014 to 2018, we study how households learn about and respond to enforcement risk by combining daily bank account transaction data with arrest-level records of Immigration and Customs Enforcement (ICE) operations. We document that enforcement follows predictable weekday patterns and that communities with large immigrant populations have learned them: Consumption is depressed not only on days when enforcement occurs, but also on the same weekday of other weeks—when no enforcement occurs—and rebounds in between, indicating that behavior tracks beliefs about enforcement risk, not just realized enforcement. Instrumenting for beliefs using the learnable structure of enforcement, we find that a 10 percentage point increase in perceived enforcement probability reduces Hispanic foreign-born consumption by approximately 5 percent. A structural model with Bayesian learning and pent-up demand reveals that roughly half of the immediate consumption decline is recovered through subsequent rebound; the rest is permanently foregone. Eliminating enforcement risk entirely would increase Hispanic foreign-born consumption by 3.6 percent, but when enforcement ceases, only 42 percent of these gains materialize within the first year —the remaining losses reflect learning frictions: even after state action stops, beliefs must still adjust.

Keywords: Immigration enforcement, Bayesian learning, beliefs, pent-up demand

JEL Codes: D83, D84, H11, J15, K37

*We especially thank Aryan Safi for his invaluable help with this project, Liam Puknys for his research assistance, and Gadi Barlevy, Kristin Butcher, Frank DiTraglia and Ezra Karger for their suggestions. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of Chicago or of the Federal Reserve System.

[†]alberto.ciancio@glasgow.ac.uk, 2 Discovery Pl, Glasgow G11 6EY, UK.

[‡]camilo.garcia-jimeno@chi.frb.org, 230 S. LaSalle Street, Chicago, IL.

1 Introduction

Governments use a variety of tools—censuses, standardized records, rule enforcement—to monitor and control populations. Populations, in turn, can respond to state action in multiple ways, creating a dynamic between state power and societal adaptation. In the Southeast Asian context, for example, James Scott has documented how populations employ everyday strategies—foot-dragging, evasion, dissimulation—to navigate state power and minimize exposure (Scott, 1985, 2009).[1]

This paper studies an informational variant in the context of U.S. immigration enforcement: we document how immigrant communities learn to anticipate Immigration and Customs Enforcement (ICE) operations, respond to *expected* enforcement (not just realized events), and adjust their daily behavior accordingly. Moreover, we quantify the consumption costs of this behavioral adaptation. By observing ICE's predictable weekday patterns of enforcement and learning from past experience, communities develop their own forecasting capacity. This adaptation partially mitigates—but does not eliminate—the costs of living under immigration enforcement risk.

Immigration enforcement affects far more than the unlawfully present individuals who are directly apprehended. By 2023, an estimated 26 million people in the United States lived in households with at least one unlawfully present immigrant (Pew Research Center, 2025). Media accounts describe them afraid to leave their homes, to go to work, or to seek medical care, and a growing academic literature documents these "chilling effects"—reductions in safety net participation, healthcare utilization, crime reporting, and labor force participation in response to enforcement activity (Alsan and Yang, 2024; Amuedo-Dorantes and Antman, 2022; East et al., 2023; Gonçalves et al., 2024; Watson, 2014). This literature typically estimates the effects of *realized* enforcement. A raid occurs, and consumption falls. Deportations increase, and labor supply declines. The underlying behavioral mechanism, however, is rarely the shock itself—it is how the shock affects *beliefs* about enforcement risk. People reduce consumption not because a raid occurred, but because a raid signals elevated future risk.

This observation motivates four questions. First, do unexpected changes in immigration enforcement affect economic behavior? This is the standard question in the chilling-effects literature, which studies how exogenous increases in enforcement intensity influence outcomes such as labor supply or program participation. Second, is immigration enforcement learnable? If ICE operates unpredictably, communities cannot form meaningful expecta-

[1] Acemoglu and Robinson (2019) go further, arguing that the competition dynamics between civil society and the state shape institutional development paths.

tions. But if enforcement follows systematic patterns—across weekdays, locations, or over time—learning becomes possible. Third, do communities learn from experience and respond to expected enforcement risk? If individuals track enforcement patterns, behavior should adjust not only when raids occur, but also on days when raids are more likely to occur, even if no raid ultimately takes place. Finally, what are the welfare costs of immigration enforcement risk, and how does learning shape these costs when enforcement patterns change?

We address these questions in the context of routine federal immigration enforcement between 2014 and 2018, using two novel data sources: daily, transaction-level financial records from debit, credit, and prepaid cards, and arrest-by-arrest records of ICE interior enforcement obtained through FOIA litigation. The financial data allow us to measure consumption activity at the ZIP Code Tabulation Area (ZCTA) by day level, while the arrest records provide precise timing and location of enforcement events.

We begin with Operation Crosscheck—a coordinated, five-day enforcement surge in March 2015 that represented a sharp deviation from routine ICE activity. Because Crosscheck was not part of ICE's usual enforcement patterns, it constitutes an unexpected enforcement shock. We find that high–Hispanic foreign-born ZCTAs experienced large and immediate declines in consumption during Crosscheck, concentrated among populations most exposed to enforcement risk. This result establishes that immigration enforcement has economically meaningful effects on consumption and confirms the presence of chilling effects following unexpected enforcement.

Having established that enforcement matters, we next ask whether enforcement is learnable. We show that ICE enforcement exhibits highly predictable patterns over time, with strong weekday effects and serial correlation. Once we condition on the day of the week and the previous day's enforcement activity, there is little remaining predictable structure. This implies that communities observing enforcement over time can form accurate forecasts of enforcement risk, creating scope for learning and expectation-based responses.

We then turn to our central question: whether communities respond to expected enforcement risk rather than to realized enforcement alone. We design event studies centered on "clean" raid events—days when a raid occurred with no other raids in the surrounding ± 8-day window. If communities learn enforcement patterns, consumption should decline not only on the raid day itself, but also on other high-risk days within the event window. Consumption is significantly depressed on the same weekday one week before and after the raid—even though, by construction, no raid occurred on those days—while it rises on low-risk weekdays within the window. These patterns provide direct evidence that behavior responds to beliefs about enforcement risk, not merely to realized raids.

The event studies demonstrate that behavior responds to expected enforcement risk. But

quantifying the causal effect of enforcement risk perceptions on economic activity requires addressing measurement error: realized raids imperfectly proxy for latent beliefs. With this purpose, we develop an instrumental variables strategy that exploits the learnable structure of enforcement: lagged same-weekday raids update beliefs without directly affecting current activity. The IV estimates reveal substantial attenuation in OLS and imply that a 10 percentage point increase in the perceived raid probability reduces Hispanic foreign-born economic activity by approximately 5 percent.

The weekday patterns we document—depressed activity on high-risk days, elevated activity on low-risk days—raise an identification question: is this consumption permanently foregone, or shifted inter-temporally to safer days? This distinction is crucial for welfare analysis, as reduced-form estimates cannot separately identify these margins. At the same time, enforcement patterns are not perfectly stable: relying on a structural break detection exercise we find 'enforcement regime' changes in roughly 30 percent of ZCTAs, primarily intensity shifts. This instability motivates a model where economic agents can adapt their learning when forecasts become unreliable.

We therefore estimate a structural model that incorporates Bayesian learning and pent-up demand incentives. In the model, suppressed activity can accumulate and partially rebound on safer days, and beliefs update with experience and can be discounted when past forecasts prove inaccurate. The model allows us to quantify the welfare costs of enforcement risk, distinguishing shifted from foregone consumption, and accounting for how learning frictions shape welfare when enforcement patterns change.

Three findings emerge. First, pent-up demand is significant: suppressed consumption generates above-baseline activity within 1–2 days (half-life of approximately 1.3 days). Second, we confirm the importance of chilling effects: higher perceived risk reduces activity. Third, individuals appear to discount unreliable beliefs, though this estimate is less precise. Impulse responses reveal roughly half the immediate chilling effect is recovered through bounce-back while half is genuinely foregone.

The estimated model enables counterfactual exercises that incorporate the full learning and pent-up demand dynamics, allowing us to quantify the value of information, the costs of learning frictions, and some welfare consequences of alternative enforcement policies. These exercises exploit the model's key features: beliefs evolve endogenously as consumers observe enforcement signals, and pent-up demand accumulates when consumption is suppressed. Comparative statics that assume instantaneous adjustment would miss the welfare costs that arise during realistic transitions.

We first quantify willingness to pay for information about enforcement. Under perfect foresight—where immigrants know exactly when raids will occur—average yearly consump-

tion activity *falls* by 1.7 percent as immigrants optimally time their behavior, staying home precisely when enforcement occurs. This consumption sacrifice can be interpreted as a measure of communities' willingness to pay for perfect information. There is considerable heterogeneity: communities with accurate forecasts would sacrifice only 0.2 percent, while those with poor forecasts would pay 3.0 percent. Separately, the willingness to pay to eliminate enforcement risk entirely—comparing a world with no enforcement to one with known average risk—is 3.6 percent of Hispanic foreign-born consumption.

We also examine alternative enforcement policies. When we simulate reducing or eliminating raids, immigrants must *learn* that enforcement has changed, starting from their historical beliefs. Full elimination yields average yearly consumption gains of 1.5 percent—roughly 42 percent of the 3.6 percent potential willingness to pay. The remaining 58 percent represents the cost of learning frictions: even after enforcement stops, beliefs take time to adjust. This finding has direct relevance for welfare evaluation: the benefits of reducing enforcement depend critically on how quickly communities learn that the environment has changed. When we instead simulate increases in enforcement intensity, we find that doubling the observed average enforcement risk leads to a 4 percent fall in economic activity among the Hispanic foreign born. Finally, we use the model to study the consequences of adaptive learning itself. Comparing economic activity under our estimated model to a counterfactual in which immigrants respond to beliefs without modulating their response as a function of cumulative forecast error, we find that this form of adaptive behavior raises economic activity by an average of about 1.8 percent among Hispanic foreign-born households, with substantially larger gains in locations where enforcement patterns are harder to predict.

Our paper makes several contributions: First, we suggest a novel interpretation of community learning about governmental action as a form of informational adaptation to navigate state power. We cleanly separate responses to *expected* enforcement (event studies with the day ± 7 effect) from responses to *unexpected* enforcement (Operation Crosscheck), showing that individuals react to novel information, and can anticipate enforcement through learning. We quantify short-run 'chilling' effects on economic activity of enforcement risk perceptions. We additionally provide suggestive evidence of learning sophistication: economic agents adjust the importance they give to their learned beliefs when forecasts prove unreliable. Finally, our framework allows us to distinguish consumption that is merely *shifted* (timing distortion) from consumption that is permanently *foregone*, a distinction that is important for welfare analysis and one that has proven elusive to quantify in previous empirical work. It also allows us to quantify the value of information about enforcement risk, and to measure the dynamic implications of several immigration enforcement policies.

Our study contributes to several areas of research. First, as mentioned above, a growing

strand of literature documents that immigration enforcement activity generates "chilling effects" on both unauthorized and lawfully present immigrants, including U.S. citizens (Watson and Thompson, 2022). We contribute to this literature by documenting chilling effects along another margin—consumption—and, importantly, by isolating immigrants' *beliefs* about enforcement risk and estimating the causal impact of beliefs on behavior.

Our paper also relates to the literature studying the economic costs of living under the threat of state action. Studies of conflict (Tapsoba (2023); Besley and Mueller (2012)), terrorism (Alfano and Görlach (2022); Abadie and Gardeazabal (2003); Draca et al. (2011)), and policy uncertainty (Pastor and Veronesi (2012)) demonstrate that expectations of future state actions depress asset values, investment, and human capital accumulation, even in the absence of direct exposure. In particular, in the economics of crime literature, deterrence is determined by the perceived probability of punishment (Chalfin and McCrary, 2017). While this literature highlights the importance of perceived risk, beliefs are typically proxied by realized events or latent risk estimated by the researcher. An exception is Lochner (2007), who uses self-reported arrest probabilities from the NLSY97 to show that deterrence depends on the perceived likelihood of arrest and that individuals update these beliefs based on their own experiences. Closely related is García-Jimeno (2016), who studies learning about law enforcement under alcohol prohibition and shows that beliefs about enforcement effectiveness play a central role in shaping compliance and political outcomes. We contribute by explicitly modeling belief formation and learning about enforcement risk and by isolating the effect of these beliefs, rather than realized enforcement, on consumption.

We adopt an adaptive learning framework inspired by Evans and Honkapohja (2001) and Orphanides and Williams (2005), where agents learn about inflation dynamics or policy rules in macroeconomic models. Adaptive learning models were first estimated in macroeconomic settings by Milani (2007) in a DSGE framework, while Slobodyan and Wouters (2012) introduced constant-gain learning, which allows the speed of learning to vary over time—an important feature when beliefs respond strongly to salient events. More recently, Carvalho et al. (2023) show that agents can recognize the un-anchoring of inflation expectations through sophisticated forms of adaptive learning. Similarly, using microdata, Malmendier and Nagel (2016) show that individuals learn about inflation through experience, with more recent experiences receiving greater weight. We contribute to the adaptive learning literature by studying belief formation in a setting where agents learn about a latent enforcement regime from realized enforcement events rather than from aggregate macroeconomic outcomes. Unlike most of this literature, we estimate beliefs using micro-level consumption behavior and high-frequency enforcement data.

2 Data and Institutional Context

2.1 Immigration Enforcement in the United States

About 5 percent of the population residing in the U.S. is unlawfully present, having entered without inspection or having overstayed their visas. Most participate in the labor force (5.6 percent of the workforce in 2023). Three quarters originate from Latin America and concentrate in California, Texas, and Florida, although this concentration has declined over time (Pew Research Center, 2025).

Immigration enforcement operates through two channels: border enforcement by Customs and Border Protection (CBP), and interior enforcement by Immigration and Customs Enforcement (ICE). We focus on interior enforcement—arrests made within communities rather than at or near ports of entry. ICE operates through 24 regional Areas of Responsibility (AORs), each covering multiple states. Because the federal government has broad discretion over enforcement, intensity varies across AORs and over time—variation central to our empirical approach.

Between 2014 and 2018, the period we study, three quarters of interior enforcement took place when ICE apprehended and transferred for deportation individuals already detained by local law enforcement agencies.[2] ICE, however, can also arrest individuals directly in the community—the focus of this paper. Though fewer in number, these arrests are highly visible and disruptive. So-called immigration raids—large-scale operations arresting hundreds simultaneously, such as Operation Crosscheck in March 2015—have received extensive media coverage and generated widespread fear. Raids take place in workplaces, parking lots, and outside homes and immigration courts.[3] Numerous media accounts describe immigrants as afraid to leave their homes.[4] In practice, most ICE arrests do not come from large operations but from daily local targeted operations with one or few targets for deportation[5]. Through-

[2]See TRAC (https://tracreports.org/phptools/immigration/arrest/). This process is overseen by the Enforcement and Removal Operations (ERO) program, which includes the Secure Communities program. Under Secure Communities, fingerprints of inmates are shared with ICE to identify potential immigration violations. Cooperation by local law enforcement agencies is not mandatory and has been the subject of political conflict (Ciancio and García-Jimeno, 2024).

[3]ICE officers derive their civil arrest authority from the Immigration and Nationality Act, which authorizes immigration officers to arrest noncitizens with or without a warrant under specified conditions (INA §§ 236, 287; 8 U.S.C. §§ 1226, 1357). Absent consent or exigent circumstances, officers generally require a judicial warrant to enter a private residence.

[4]For example, "Immigrant Communities in Hiding: 'People Think ICE Is Everywhere,'" *The New York Times*, January 30, 2025, and "ICE raids are looming. Panicked immigrants are skipping work, hiding out and bracing for the worst," CNN, July 12, 2019.

[5]These targeted operations sometimes lead to collateral arrests, arrests of individuals who were not the original targets of an operation but are encountered during its execution and determined to be candidates for removal by the officers in charge.

out the paper, we use the term "raid" to describe community-based ICE arrest operations, including those involving a single target.

We study the period from October 2014 through May 2018, straddling two administrations with different enforcement priorities. The Obama administration emphasized enforcement against those with serious criminal convictions; the Trump administration expanded enforcement efforts to a broader population. This policy shift, combined with daily data, generates variation in both the level and predictability of enforcement. Figure 1 displays weekly interior arrests over the sample period, with vertical lines marking key policy events.

Perceived Enforcement Risk. Throughout the paper we focus on *perceived enforcement risk*: the expectation, held community members, that immigration enforcement will occur in their location on a given day. This perceived risk is not directly observed by the researcher. It may be shaped by learning from past enforcement activity, by information transmitted through social, organizational, or institutional channels, or by a combination of both. Our analysis studies how economic behavior responds to variation in perceived enforcement risk, and how that risk evolves with observed enforcement-related information over time.

2.2 Data

Immigration Enforcement. Data on ICE arrests come from the Transactional Records Access Clearinghouse (TRAC) at Syracuse University, which obtained arrest-by-arrest records through FOIA litigation. We use arrests classified as "community (at large)"—those outside jails or prisons. Each record includes the arrest date, county[6], and most serious criminal conviction (categorized as level 1, 2, or 3, or none). Around 25% of arrests fall into the top priority (level 1) category and a further 10% fall in the second most serious category. The most common categories are no conviction (28%), driving under influence (16%), traffic offense (5%), assault (4%), drug possession (4%) and illegal re-entry (4%).

Because our consumption data are at the ZIP Code Tabulation Area (ZCTA) level, we map county arrests to ZCTAs using Census relationship files. ZCTAs can span multiple counties; we associate a ZCTA with each county containing at least 2.5 percent of its population, then sum arrests across associated counties.[7] The resulting data include 17,765 ZCTAs with at least one arrest during the sample period. Figure A1 displays the geographic distribution of arrests.

[6]ICE does not officially track arrest locations; TRAC reconstructed county information through hundreds of FOIA requests and cross-referencing multiple sources. In some cases, the county designation includes surrounding areas where exact location could not be isolated.

[7]We provide details of this mapping procedure in Appendix A.

Economic Activity. We measure economic activity using data from Facteus, a financial data aggregator that provides transaction-level records from a very large sample of debit, credit, and prepaid cards. We work with the subset of accounts with at least one transaction before the beginning of, and after the end of our sample period to ensure our analysis holds constant the set of accounts of over time, and we aggregate these data to the ZCTA level.

Our main outcome is the *share of accounts active*—accounts with at least one transaction in the ZCTA-day, divided by total accounts. This captures the extensive margin: whether individuals transact at all, rather than how much they spend. The extensive margin suits avoidance behavior: if people stay home, they make fewer transactions. In robustness exercises we also use the total monetary value of the transactions. The balanced panel contains approximately 30,000 ZCTAs observed daily. We provide additional details on these data and the construction of our sample in Appendix A.1.

Demographics. ZCTA-level characteristics come from the 5-Year 2018 American Community Survey. Our key variable is the *Hispanic foreign-born share*—the fraction both Hispanic and foreign-born—which proxies for the population most exposed to enforcement risk. We also use median household income, total population, and racial composition.

2.3 Sample Construction and Summary Statistics

The final sample merges the three data sources at the ZCTA-date level. We impose a "selection cone" requiring at least 30 accounts and 3–25 accounts per 1,000 population, ensuring adequate coverage while excluding implausible penetration rates. This yields 12,865 ZCTAs over 1,339 days (17.2 million observations). Of these, 73.8% experienced at least one arrest ("ever-raided"). Table 1 presents summary statistics split by median Hispanic foreign-born share (1.7%). The share of accounts active averages 10.5%, similar across groups. High-Hispanic-foreign-born ZCTAs experience more enforcement: 18.4% of days have at least one raid versus 6.7%, and 84.1% are ever-raided versus 63.5%.

3 Beliefs and Patterns of Immigration Enforcement

3.1 Setup

A *chilling effect* is the causal response of economic activity to changes in beliefs about enforcement risk—not to enforcement itself. Absent variation in risk perceptions, chilling effects cannot be measured. But even when beliefs change—from news, policy announcements, or observed enforcement—they remain unobserved.

Studies of chilling effects take one of two approaches: (i) exploit surprise enforcement announcements, assuming these shift beliefs, or (ii) use observed enforcement as a proxy for beliefs. Either way, to detect chilling effects learning must occur. Conversely, estimating chilling effects provides evidence of belief updating.

Consider a linear model of economic activity:

$$y_{z,t} = \alpha + \beta r^e_{z,t} + \epsilon_{z,t} \tag{1}$$

where $r^e_{z,t} \equiv \mathbb{E}[r_{z,t}|\mathcal{I}_{z,t}]$ is the subjective probability of a raid in community z on day t, $\mathcal{I}_{z,t}$ is information available at the beginning of the period, and β captures the chilling effect. We assume $\mathbb{E}[\epsilon_{z,t} r_{z,t} \mid \mathbf{x}_{z,t}] = 0$ for an appropriate set of controls $\mathbf{x}_{z,t}$—ICE's enforcement decisions do not respond to daily economic shocks. Because we aggregate data at the ZCTA level, we also assume beliefs are common within communities. If we observed beliefs, regressing activity on them with ZCTA and time fixed effects would recover β.

3.2 Operation Crosscheck

We begin with Operation Crosscheck, a coordinated enforcement surge that provides clean identification of responses to *unexpected* enforcement. This establishes a baseline chilling effect: enforcement reduces subsequent economic activity in locations with large numbers of immigrants. In subsequent sections, we show that ICE enforcement follows predictable patterns that communities have learned, responding not just to realized but to *expected* raids.

3.2.1 Background

Operation Crosscheck was a coordinated five-day enforcement action by ICE from March 1–5, 2015, resulting in 2,059 arrests nationwide.[8] Unlike routine enforcement, this was a planned surge executed simultaneously across all ICE districts.

Crosscheck as an Information Shock. Crosscheck is valuable as a natural experiment because it represents *unexpected* enforcement. Three features distinguish it from routine

[8] See https://www.dhs.gov/news/2015/03/09/2059-criminals-arrested-ice-nationwide-operation. Operation Cross Check is part of ERO's National Fugitive Operations Program and is described as targeting specific "priority" populations (e.g., at-large criminals with violent convictions or transnational gang members). Officers "crosscheck" identity and location leads across immigration enforcement records and law-enforcement/criminal-history information to produce arrestable targets, then execute coordinated arrests nationwide. Previous Operations Crosscheck took place before the beginning of our sample period. A second Operation Crosscheck took place in January 2017. Because its impact overlaps with President Trump's first inauguration, we do not use it here.

enforcement. First, the operation was planned nationally and executed simultaneously—timing was not predictable from local conditions. Second, Crosscheck began on a Sunday and ran through Thursday, deviating from the typical weekday concentration we will document in Section 3.3. Third, the scale was unprecedented: while typical enforcement involves 1.2 arrests, Crosscheck-affected ZCTAs saw dozens simultaneously. Because communities could not have anticipated Crosscheck from past patterns, any response represents reaction to unexpected news rather than pre-planned avoidance.

3.2.2 Event Study Design

Let event time dummies be

$$D_{z,t}^{k} = \mathbf{1}[\tau_{z,t} = k], \qquad k \in \{-K, \ldots, +K\}, \tag{2}$$

and consider the baseline event study specification

$$y_{z,t} = \alpha_z + \delta_t + \sum_{k \neq 0} \gamma_k D_{z,t}^{k} + \varepsilon_{z,t} \tag{3}$$

where $y_{z,t}$ is economic activity, α_z are ZCTA fixed effects, and δ_t are date fixed effects. The coefficients γ_k capture mean differences in the outcome k days from baseline $k = 0$. In practice, our focus is on whether economic activity changes differentially in communities with larger fractions of people subject to immigration enforcement risk, so our specification allows for heterogeneity along the Hispanic foreign born margin, h_z: $\gamma_k = \gamma_k^0 + \gamma_k^1 h_z$.

Crosscheck represents a surprise signal S_τ that expands the information set to $\mathcal{I}_{zt} \cup S_\tau$ for periods $\tau > 0$, inducing a belief innovation. Denoting with * variables after partialling out fixed effects, the least squares estimator identifies

$$\gamma_k = \beta \underbrace{\left(\mathbb{E}\left[r_{zt}^{e*}(\mathcal{I}_{z,t} \cup S_\tau) \mid \tau = k\right] - \mathbb{E}\left[r_{zt}^{e*}(\mathcal{I}_{z,t}) \mid \tau = 0\right]\right)}_{\text{Belief innovation from surprise}} + \underbrace{\left(\mathbb{E}\left[\epsilon_{zt}^{*} \mid \tau = k\right] - \mathbb{E}\left[\epsilon_{zt}^{*} \mid \tau = 0\right]\right)}_{\text{Residual difference}} \tag{4}$$

where expectations are over ZCTAs, with τ indexing event time. Under the assumption that Crosscheck timing is exogenous to local consumption shocks, the residual difference is zero, and the coefficients γ_τ recover the chilling effect β scaled by the average belief innovation. Before the signal arrives, $S_\tau = \varnothing$, so we expect flat pre-trends followed by a dip when the information set expands.

3.2.3 Empirical Specification

We construct an event study sample with clean windows. We include ZCTAs that (i) experienced at least one arrest during Crosscheck (March 1–5, 2015), and (ii) had no arrests in the 10 days before or after. The second requirement ensures that other enforcement actions do not contaminate the event study windows. We measure event time relative to Tuesday, March 3—the middle of the Crosscheck week. About 30 percent of ZCTAs experienced multiple arrests during the week, so we center on the midpoint. This yields 802 ZCTAs with clean Crosscheck events.

We use the two-stage approach of Gardner et al. (2024) for heterogeneous treatment effects in staggered designs. In the first stage, we estimate ZCTA-by-weekday and date fixed effects using untreated observations (ZCTA-days without raids in a ten-day window):

$$y_{zt} = \alpha_{z \times d(t)} + \delta_t + \varepsilon_{zt} \tag{5}$$

where $d(t) \in \{\text{Mon}, \ldots, \text{Sun}\}$ denotes the weekday. The ZCTA-by-weekday effects absorb systematic day-of-week patterns, ensuring we identify the Crosscheck effect from the enforcement shock rather than regular weekly cycles. Because we estimate them on untreated observations only, treatment does not pollute these fixed effects. And because we have a very large sample of untreated ZCTAs and days, we can estimate them precisely.

In the second stage, we regress the residualized outcome $\tilde{y}_{zt} = y_{zt} - \hat{\alpha}_{z \times d(t)} - \hat{\delta}_t$ on the interactions between the Hispanic foreign-born share h_z and event time indicators, or more flexibly, with dummies for the deciles of the Hispanic foreign born share distribution. We base our inference on 1,000 bootstrap replications clustered at the ZCTA level.

3.2.4 Identification

The ZCTA-by-weekday fixed effects absorb both regular consumption patterns and the predictable component of enforcement documented in Section 3.3. What remains is the *surprise*: Crosscheck's deviation from expected patterns. We therefore need not model beliefs explicitly at this stage.

3.2.5 Results

Figure 2 presents coefficients on the interaction terms on the event-time by Hispanic foreign-born share. The coefficients show a sharp negative dip on Monday and Tuesday (event time -1 and 0)—the core of the operation—indicating that ZCTAs with higher Hispanic foreign-born shares experienced reduced economic activity during the enforcement surge. The effect

is sizable: for the average ZCTA in the Crosscheck sample, consumption drops by about 8% of baseline activity during the peak of the operation. Communities with larger at-risk populations curtailed activity in response to unexpected enforcement.

Figure A4 presents results by Hispanic foreign-born share decile. The negative effect on Monday and Tuesday concentrates in top deciles: deciles 1–3 show clear, statistically significant dips, while deciles 5–9 remain flat around zero. This suggests fear among at-risk populations, rather than some channel affecting all communities equally.

The Crosscheck analysis establishes a familiar but important baseline: unexpected enforcement reduces immediately subsequent economic activity in immigrant communities. But if communities respond to unexpected enforcement, do they also respond to *expected* enforcement? Have they learned ICE patterns well enough to adjust behavior in anticipation of raids? Answering this requires understanding whether ICE enforcement is predictable. In the next section, we show that it is: raids follow strong weekday patterns and serial correlation that communities could learn. We then show that consumption tracks these expected patterns—not just realized raids.

3.3 ICE Enforcement Patterns

3.3.1 Weekday Structure

ICE enforcement follows pronounced day-of-the-week patterns. Figure 3a displays the partial autocorrelation function (PACF) of the raid indicator, showing strong autocorrelation at lags that are multiples of seven days. ICE operations concentrate on weekdays rather than weekends, and these patterns persist within ZCTAs over time.

These weekday patterns vary across ZCTAs. Figure 4 displays a heatmap of weekday enforcement profiles, with each row representing one ZCTA and columns showing the share of raid-days on each weekday. We group ZCTAs by peak enforcement day and sort them by concentration within each group. While mid-week enforcement is generally higher (visible as the vertical structure), intensity varies across ZCTAs (visible as the horizontal variation). Some ZCTAs have concentrated enforcement on a single day; others have diffuse profiles. This suggests communities cannot simply adopt the aggregate pattern, and must learn their own local enforcement rhythms.

The weekday structure likely reflects operational constraints: ICE officers work regular schedules, courts and processing facilities operate on weekday hours, and multi-day operations may be planned around the work week. Although it is difficult to establish the underlying reasons for these enforcement regularities, the result is a predictable pattern.

3.3.2 Serial Correlation

Beyond weekday patterns, enforcement exhibits strong serial correlation. A raid today increases the probability of a raid tomorrow, even controlling for weekday. This likely reflects multi-day operations—when ICE targets an area, officers conduct arrests over several days before moving on. Conditional on weekday, yesterday's raid is the strongest predictor of today's. This creates an additional learnable signal. Both the weekly signature and the lag-1 dependence will inform our modeling of belief updating in Section 4.3.

3.3.3 Is ICE Enforcement Learnable?: Residual Autocorrelation

Can individuals grasp predictable patterns using a simple forecasting rule? Figure 3b displays the partial autocorrelation function (PACF) of the raid indicator—after regressing it on weekday dummies and the lagged raid indicator. The autocorrelation is flat: once we control for weekday and the first lag, no predictable structure remains. Thus, two features characterize ICE enforcement: (1) weekday patterns varying across ZCTAs, and (2) serial correlation captured by the first lag. A learner conditioning on weekday and yesterday's outcome can form accurate forecasts, as long as patterns remain stable. These findings open the room for the next section's question: do communities actually learn these patterns? If so, consumption should track *expected* raid likelihood, not just realized raids.

3.4 Event Studies: Inferring Beliefs about Enforcement Risk

Do communities actually learn ICE's weekday patterns? We answer this with event studies centered on "clean" raid events: days when a raid occurred with no other raids in the surrounding ± 8 day window. By centering on realized raids, we anchor on days when enforcement risk was high (since raids cluster on preferred weekdays). If communities track ICE patterns, consumption should dip not only on the raid day itself, but on all high-likelihood days within the window—including days -7 and $+7$, where by construction no raid occurred.

3.4.1 Event Study Design

The event-time dummies

$$D^k_{z,t} = \mathbf{1}[\tau_{z,t} = k], \qquad k \in \{-8, \ldots, +8\}$$

are now defined relative to a day $\tau = 0$ when $r_{z,\tau} = 1$, and the baseline specification is the same as in (3). The coefficients γ_k now capture the average difference in economic activity

k days away from a 'high enforcement risk' day, relative to the baseline day.

The least squares estimator of the $\boldsymbol{\gamma} = \{\gamma_k\}$'s from (3) now identify

$$\gamma_k = \beta \underbrace{\left(\mathbb{E}[r^{e*}_{z,t} \mid \tau = k] - \mathbb{E}[r^{e*}_{z,t} \mid \tau = 0]\right)}_{\text{Belief gradient in event time}} + \underbrace{\left(\mathbb{E}[\epsilon^{*}_{z,t} \mid \tau = k] - \mathbb{E}[\epsilon^{*}_{z,t} \mid \tau = 0]\right)}_{\text{Residual}} \tag{6}$$

where the expectations are taken over the cross-section of ZCTA-date pairs, with τ indexing event time. The key insight is that by centering event windows on realized raids ($r_{z,t} = 1$ at $\tau = 0$), we align clocks across ZCTAs at times when beliefs about enforcement risk are plausibly elevated. Because ICE follows weekday patterns, day 0 is likely a high-probability weekday for that ZCTA, making days ± 7 (the same weekday) also high-probability, and different weekdays, say ± 4, lower-probability. This converts idiosyncratic weekday patterns into a common event-time gradient that we can estimate.

Without anchoring on realizations, there would be no systematic reason for beliefs to vary with event time after including ZCTA and date fixed effects. Anchoring on realizations generates the identifying variation. While Crosscheck recovers responses to belief innovations from surprises, the clean raid event studies here recover the cross-weekday belief gradient.

3.4.2 Temporal Alignment between Raid Risk and Economic Activity

As with Crosscheck, we use "clean" windows where a raid occurred at $\tau = 0$ with no raids in $[-8, +8]$. Our sample contains 33,091 clean raid events across 5,967 ZCTAs (562,547 ZCTA-day observations; Appendix Figure A5 shows the geographic distribution). We implement the Gardner et al. (2024) two-stage procedure, allowing γ_k to vary with the Hispanic foreign-born share.

Figure 6 presents our central finding. Orange points show consumption coefficients along the Hispanic foreign-born share gradient. The blue line overlays the average raid likelihood by event day from an autocorrelation regression (see Appendix Figure A6). The correspondence is striking: consumption closely mirrors ZCTA-level enforcement risk. When the raid likelihood is high (days -7, 0, $+7$), consumption is depressed; when the likelihood is low (days -4, $+4$), consumption recovers.

The event-study patterns indicate that economic activity responds to *perceived enforcement risk* rather than to realized enforcement alone. Consumption is depressed not only on the day of a raid, but also on days that historically carry a higher likelihood of enforcement—most notably the same weekday one week before and after the raid—even though no enforcement occurs on those days by construction. This alignment between activity and historically high-risk weekdays suggests that households condition their behavior on expec-

tations about enforcement risk associated with particular days, rather than reacting only to contemporaneous enforcement events.

These anticipatory responses are consistent with belief formation based on past enforcement experience, but they do not require a specific micro-level mechanism. Perceived enforcement risk may reflect learning from prior enforcement realizations, information transmitted through social or organizational networks, institutional coordination (e.g., employer scheduling), or a combination of these channels. What we identify in these exercises is that behavior responds systematically to variation in perceived enforcement risk across days.

3.4.3 Robustness

We also estimate the belief gradient by Hispanic foreign-born share decile (Figure A7). The seven-day pattern is strongest in the highest deciles, attenuates in deciles 4–6, and disappears in the bottom deciles. Effects concentrate exactly where unlawfully present populations are likely to be the largest, supporting our interpretation. For this reason, we restrict our subsequent analysis (IV and structural estimation) to ZCTAs in the top quintile of the Hispanic foreign-born share.

Another concern about our results there is that the weekly pattern could reflect weekday-specific consumption (e.g., people shop more on Saturdays) rather than risk perceptions. Date fixed effects absorb aggregate weekday effects but not ZCTA-specific patterns. Formally, if $\epsilon_{zt} = \theta_z^{d(t)} + u_{zt}$ where $\theta_z^{d(t)}$ captures ZCTA-specific weekday effects, the residual in (6) becomes $\mathbb{E}[\theta_z^k] - \mathbb{E}[\theta_z^{d(0)}]$. Since $k = \pm 7$ share the same weekday as $k = 0$, we have $\mathbb{E}[\theta^{\pm 7}] = \mathbb{E}[\theta^0]$. But for other weekdays, $\mathbb{E}[\theta_z^k]$ could vary, generating a seven-day pattern even with $\beta = 0$.

We address this by replacing ZCTA fixed effects with ZCTA $\times$ weekday fixed effects, absorbing $\theta_z^{d(t)}$. Appendix Figure A8 shows the overlay plot with these fixed effects. The pattern survives, though attenuated—as expected, since weekday fixed effects absorb some belief-driven variation (beliefs themselves have weekday structure). The pattern is also robust to using log spending as the outcome (Appendix Figure A9).

3.4.4 From Beliefs to Chilling Effect Magnitudes

The event study shows that consumption tracks expected raid risk, rather than reacting only to realized enforcement. But the coefficients γ_k are proportional to both the belief gradient and the chilling effect β; they do not directly identify β. Taken together, the evidence establishes three facts: (i) communities respond strongly to unexpected enforcement (Crosscheck), (ii) ICE enforcement follows predictable temporal patterns, and (iii) economic

activity varies systematically with high- and low-risk days. Together, these motivate an identification strategy that exploits predictable enforcement-related information to recover average chilling effects, which we present next.

3.5 Measurement Error in the Estimation of Chilling Effects

The event studies in Section 3.4 document that consumption tracks expected raid likelihood—communities with relatively large Hispanic foreign-born communities show significant knowledge about ICE's enforcement patterns. Moreover, the exercise leveraged the idea that a day with an immigration enforcement action is likely a day when risk perceptions are high. That is, we used observed raids as proxies for beliefs about enforcement risk. Here we bring together these two ideas to show how to recover an average chilling effect when there is learning about immigration enforcement patterns.

3.5.1 Realizations as Proxies for Beliefs

Studies of enforcement effects typically estimate responses to enforcement shocks. But behavior responds to *beliefs* about risk, not shocks themselves, creating a measurement error problem: the realized shock is a noisy proxy for the latent belief.[9] To formalize this, recall from equation (1) that economic activity responds to beliefs if $\beta \neq 0$. Because beliefs are unobserved, a natural approach is to use the realized raid indicator r_{zt} as a proxy for r^e_{zt}. Define the prediction error as $\tilde{\eta}_{zt} \equiv r_{zt} - r^e_{zt}$—the gap between the raid realization and what agents expected. Substituting $r^e_{zt} = r_{zt} - \tilde{\eta}_{zt}$ into the model yields:

$$y_{zt} = \alpha_z + \delta_t + \beta r_{zt} + (\epsilon_{zt} - \beta\tilde{\eta}_{zt}) \tag{7}$$

The composite error now includes the prediction error $\tilde{\eta}_{zt}$, which is mechanically correlated with the regressor r_{zt}.

Measurement Error Bias. Under the maintained assumption that $\text{cov}(r_{zt}, \epsilon_{zt}) = 0$ conditional on fixed effects, the OLS estimand from (7) takes the form (see Appendix D for derivation):

$$\beta_{OLS} = \beta \cdot [\mathbb{E}[r^e_{zt}|r_{zt} = 1] - \mathbb{E}[r^e_{zt}|r_{zt} = 0]] \tag{8}$$

[9] In other contexts—such as the effect of interest rate changes on consumption—both the shock itself and beliefs about future shocks affect outcomes, creating an inherent identification challenge. Here, the consumption decision must be made before observing whether a raid occurs, so only beliefs matter at decision time, a convenient feature of our setting.

The bias factor in brackets measures how much beliefs differ between raid days and non-raid days. Three cases clarify how this bias depends on belief quality:

1. **Uninformative beliefs.** If r^e_{zt} is uncorrelated with r_{zt}, the bias factor is zero and OLS identifies nothing. This is unsurprising: using the enforcement activity as a proxy is useless when people's beliefs contain no information about the underlying enforcement activity.

2. **Oracle beliefs.** If agents perfectly forecast raids ($r^e_{zt} = r_{zt}$), the bias factor equals one and OLS is unbiased.

3. **Partial learning.** When beliefs are informative but imperfect—the empirically relevant case—the bias lies between these extremes. Moreover, OLS recovers the correct sign as long as beliefs contain some information: formally, whenever $\text{cov}(r^e_{zt}, \mathbb{P}(r_{zt} = 1|r^e_{zt})) \geq 0$ (see Appendix D).

3.5.2 Using Learning to Construct Instruments

Learning suggests a solution: if individuals form beliefs from past enforcement, lagged enforcement—particularly on the same weekday—should shift beliefs without directly affecting current consumption. We construct an instrument as follows. For each ZCTA-weekday-date, we compute the rolling mean of raid indicators over the previous ℓ weeks *on the same weekday*. The identifying assumption underlying this instrumental variables strategy is that lagged same-weekday enforcement affects current economic activity only through its effect on *perceived enforcement risk*. Because individuals must decide whether to go out/consume before observing enforcement on that day, contemporaneous behavior responds to expectations about enforcement risk rather than to realized raids themselves. Recent enforcement events provide salient information about the likelihood of enforcement on similar days, shifting perceived risk even if no raid ultimately occurs. Conditional on location and weekday fixed effects, lagged same-weekday enforcement therefore shifts behavior through changes in perceived enforcement risk, rather than through direct effects on contemporaneous economic activity.

Our IV strategy relies on the idea that agents form beliefs r^e_{zt} using past information, so lagged realizations of the signal r_{zt-k} belong to $\mathcal{I}_{z,t-1}$ and predict r^e_{zt}. It does not require informationally optimal beliefs (i.e., $r^e_{zt} = \mathbb{E}[r_{zt}|\mathcal{I}_{z,t}]$). Instead, three high-level conditions are sufficient: (i) relevance: $\text{Cov}(r^e_{zt}, r_{z,t-k}) \neq 0$ for some lag k (e.g., same-weekday $k = 7$); (ii) lag-orthogonality of the belief–realization gap: $\mathbb{E}[\eta_{zt}|\mathcal{I}_{z,t-1}] = 0$; and (iii) outcome exogeneity: $\mathbb{E}[\epsilon_{zt}|\mathcal{I}_{z,t-1}] = 0$. Under (i)–(iii), the population moment $\mathbb{E}[r_{z,t-k}(y_{zt} - \beta r_{zt})] = 0$ holds, so

2SLS using $r_{z,t-k}$ identifies β. This weaker orthogonality is similar to that from applied-micro "learning from experience" designs that use lagged own/peer realizations as belief shifters (e.g., Conley and Udry (2010); Munshi (2004)), while avoiding assumptions of fully optimal Bayesian updating. In our pooled panel with date and ZCTA by weekday fixed effects, we use same-weekday lags (e.g., $t-7$, $t-14$) as instruments.

3.5.3 Instrument Selection

Figure 5 displays first-stage statistics across window lengths, ℓ. Given our findings above, we restrict the sample here to ZCTAs in the top quintile of the Hispanic foreign-born share distribution. Panel (a) shows the first-stage coefficient; Panel (b) shows the F-statistic. The first-stage relationship exhibits a clear pattern as a function of lag length window: the coefficient increases up to around 13 weeks and flattens thereafter, while the F-statistic show an inverted-U shape that peaks at 13 weeks. This pattern is consistent with belief dynamics under which recent and past enforcement actions contain relevant information, but where people place greater weight on recent signals as in learning models of adaptive gain learning. Purely static or permanent information mechanisms, or institutional responses whose salience fades over time would not generate such pattern. All specifications yield F-statistics well above conventional thresholds (the smallest exceeds 3,000), so weak instruments are not a concern. We use the 13-week window as our baseline.

3.5.4 Results

Table 2 presents IV and OLS estimates. Column (1) shows OLS with ZCTA $\times$ weekday and date fixed effects: the coefficient is negative but very small and statistically insignificant (-0.0001, s.e. $= 0.0002$). As expected: ZCTA $\times$ weekday fixed effects absorb the predictable component of beliefs, leaving only noise and driving OLS toward zero.

Column (2) presents the IV estimate using the 13-week rolling mean with the same fixed effects. The coefficient increases to -0.007, statistically significant at the 5% level. The correction is large but sensible since the realized raid indicator is a crude proxy—it equals one only when a raid occurs, even though beliefs may be elevated on any historically high-probability day.

The remaining columns probe the sensitivity of this estimate to the fixed effects structure and instrument specification. Column (4) omits ZCTA and weekday controls entirely, including only date fixed effects; the IV coefficient jumps to -0.036, much larger than our baseline. This inflation highlights the importance of accounting for cross-sectional and cross-weekday heterogeneity in enforcement exposure that the instrument otherwise picks up.

Indeed, once we include ZCTA and weekday fixed effects in any form, the estimate stabilizes. Column (3) uses additive (non-interacted) ZCTA, weekday, and date fixed effects; the coefficient is −0.012, similar to the baseline. Column (5) returns to the interacted ZCTA × weekday specification but uses an overidentified model with 1-week and 2-week same-weekday lags as separate instruments; the coefficient is −0.009, again in the same range. The consistency across columns (2), (3), and (5)—whether the ZCTA and weekday effects are interacted or additive, whether the model is exactly identified or over-identified—suggests that the key requirement is controlling for location-specific weekday patterns in some form, not the precise specification of those controls.

3.5.5 Interpretation

The IV estimates recover chilling effects operating through expectations about enforcement risk. Whether perceived risk evolves through individual learning from past enforcement, information transmission within communities, or institutional coordination does not affect the interpretation of the IV estimates as long as past enforcement activity modifies communities' information sets.

These estimates imply that the mean effect of raid risk perceptions on economic activity is large: if they increase by 10 percentage points (e.g., an increase from 0.10 to 0.20, roughly a doubling relative to the mean raid rate) in a ZCTA at the median Hispanic foreign-born share among our high-exposure sample (approximately 14 percent), the share of active accounts falls by $0.007 \times 0.1 = 0.07$ percentage points. This implies a consumption reduction of approximately 5 percent among the Hispanic foreign-born.[10]

Studies using realized enforcement shocks may substantially understate responses to enforcement *risk*. The bias is not endogeneity but measurement error: a mismatch between what we observe (shocks) and what drives behavior (beliefs). Our IV strategy relies on the underlying learning that drives enforcement risk beliefs—lagged enforcement is informative about current perceptions, and in particular the weekday structure of those risk perceptions.

3.6 Regime Changes in ICE Enforcement

How stable are ICE's enforcement patterns? If the agency occasionally changes strategy, communities face a more complex learning problem. We document that despite overall predictability, ICE does change patterns at specific times and places, creating regime changes that communities must adapt to.

[10] If the consumption fall is only among the Hispanic foreign born, then for this group it falls by $0.07/0.14 = 0.5$ percentage points. This is 5 percent of the 10 percent of daily accounts active average in the sample.

3.6.1 Detecting Regime Changes

We apply the Bai and Perron (2003) structural break detection method to weekday-specific raid probabilities, testing for significant changes in these probabilities within ZCTAs. The method estimates break dates by finding the partition of the sample that globally minimizes the sum of squared residuals, using an efficient dynamic programming algorithm.[11]

For each ZCTA, we estimate

$$r_{z,t} = \sum_{d=\text{Mon}}^{\text{Sun}} \beta_d^{(j)} \cdot \mathbf{1}[d(t)] + \gamma^j \cdot r_{z,t-1} + \epsilon_{z,t} \tag{9}$$

where $j = 0, 1, ...$ denote the underlying regimes to be detected, $r_{z,t}$ is a raid indicator (or count), $\mathbf{1}[d(t)]$ are weekday dummies, and $r_{z,t-1}$ captures serial correlation. The structural break algorithm identifies dates where the weekday coefficients $\{\beta_d\}$ shift significantly.[12]

To ensure robustness, we estimate breaks using three specifications for the dependent variable: a raid dummy, raid counts, and log raid counts. We grade detected breaks based on agreement across specifications: a break receives an "A" grade if all three specifications identify a break within 7 days of each other; a "B" grade if they agree within 28 days; and a "C" grade if at least two specifications agree within 14 days. We restrict this analysis to ZCTAs in the top quintile of the Hispanic foreign-born share —the communities we identified in our event study exercises as responding to immigration enforcement risk—, with at least 191 raid days —with sufficient enforcement activity to reliably detect pattern changes.

3.6.2 Prevalence and Timing of Regime Changes

Figure A10 displays detected break dates. We identify 513 breaks in 439 ZCTAs, and no breaks for the remaining ZCTAs. The distribution spikes in January 2017 (around President Trump's first inauguration), but the mean break date is August 31, 2016—regime changes occurred throughout the sample.

Figure A11 presents break timing by ICE district. Some districts (Boston, Los Angeles) show tight clustering around January 2017; others (San Francisco, San Antonio) are more

[11] Since partitioning the sample will always weakly reduce the SSR even absent true breaks, a formal statistical test is required. The test compares the minimized SSR under the alternative of k breaks against the null of no breaks, and uses critical values that account for the choice a maximal test statistic. To determine the number of breaks, the method evaluates whether adding an additional break significantly improves fit over a model with fewer breaks. This test requires fixing the minimum segment (regime) length as a parameter. We fix it to 15 percent of the sample (200 days), allowing for up to 6 breaks.

[12] The Bai and Perron (2003) approach allows a 'partial' regime change model where coefficients on covariates other than the main ones of interest remain constant across regimes, or a 'pure' regime change model where they are also allowed to vary. In our case, we implement the 'pure' model, allowing the coefficient on the lagged raid variable, γ^j, to also vary across regimes.

dispersed. Break dates cluster within districts, suggesting enforcement shifts occur largely at the AOR level.

3.6.3 Nature of Regime Changes: Intensity vs. Re-shuffling

An important question is whether regime changes reflect shifts in overall enforcement *intensity* (more or fewer raids on all days) or *re-shuffling* of weekday patterns (different days become high- vs. low-activity days). The distinction matters for learning: intensity shifts are easier to detect and adapt to, while re-shuffling of patterns requires relearning which days are risky.

To assess this, we compute the correlation between weekday coefficients across regimes for each ZCTA with a detected break. If a break reflects pure intensity shift, the weekday coefficients in the new regime should be highly correlated with those in the old regime (the ranking of days by risk is preserved). If instead the break reflects re-shuffling, the correlation should be low or negative.

Figures A12a and A12b display these correlations for Trump-window breaks (within one month of the inauguration) and all other breaks. Both distributions center on large positive values: regime changes primarily reflect intensity shifts, not re-shuffling. The Trump transition intensified enforcement but did not alter which days were high or low risk. ICE does not appear to shuffle weekday patterns to maintain unpredictability.

3.6.4 Implications for Learning

The presence of regime changes has implications for how communities form beliefs about enforcement risk. If enforcement patterns were perfectly stable, a simple decreasing-gain learning rule (e.g. a Kalman filter) would eventually converge to the true raid probabilities. In such context, weighting all past observations equally is informationally optimal, and beliefs become increasingly precise over time.

But with regime changes, standard Bayesian updating would fail to converge to the underlying enforcement rates. When ICE changes its patterns, old observations become misleading—they reflect a stale enforcement regime rather than current policy. Agents face a trade-off: weighting recent observations more heavily allows faster adaptation to new patterns but increases noise. This trade-off motivates the learning model specification we will introduce below, where the effective learning rate responds to forecast errors.

The finding that regime changes primarily involve intensity shifts rather than re-shuffling has a silver lining: while agents must update their beliefs about overall risk levels, they can continue to rely on previously learned weekday patterns. This reduces the learning burden

and we believe it helps explain why communities can maintain belief-tracking behavior even in an environment with occasional policy changes.

4 Model

4.1 Motivation

The reduced-form analysis in Section 3 established several key facts that motivate the model we present below:

1. **Responsiveness to enforcement risk** (Section 3.2): Communities with large fractions of Hispanic foreign born individuals reduce economic activity in response to salient enforcement news, as demonstrated by the consumption drop following the start of Operation Crosscheck.

2. **Predictable enforcement** (Section 3.3): ICE raids follow weekday patterns and exhibit serial correlation that communities can learn.

3. **Belief-tracking behavior** (Section 3.4): Consumption responds to expected raids, not just realized raids, indicating that communities have learned enforcement patterns and act on their beliefs.

4. **Detectable short-run chilling effects** (Section 3.5): In ZCTAs in the top quintile of the Hispanic foreign-born share distribution, expectations of increased enforcement intensity reduce economic activity.

5. **Regime changes** (Section 3.6): ICE occasionally shifts its enforcement intensity, which may induce individuals to be adaptive in their learning about enforcement risk.

The reduced-form event studies, however, cannot help us distinguish whether chilling effects fully explain the weekly patterns of consumption dip and surge that align with average enforcement risk, or whether *pent-up demand* partly drives this dynamic: households who avoided transactions during high-risk periods may want to catch up in later days. Whether consumption is subject to pent-up demand dynamics or not leads to different welfare implications.

The structural model below formalizes these features and allows us to distinguish between, and quantify, the importance of learning and pent-up demand in driving consumption dynamics in the presence of chilling effects from immigration enforcement risk perceptions.

We specify how agents make daily going-out decisions trading off current utility against accumulated pent-up demand. We then describe how agents form and update beliefs about enforcement risk based on observed raids. The model allows us to estimate the parameters governing i) the response of risk perceptions on economic activity —chilling effects—, ii) belief formation and adaptive responses, and iii) pent-up demand dynamics.

4.2 A Dynamic Tracking Model of Going-Out Decisions with Pent-Up Demand

4.2.1 Individual Problem

We consider a community z consisting of a continuum of individuals i of two types τ: those susceptible to immigration enforcement, H, and those not subject to immigration enforcement, N. The fraction of H types is h_z. At the beginning of every day $t = 1, 2, ...,$ each person decides whether to 'go out' ($a_{izt} = 1$) or 'stay home' ($a_{izt} = 0$). Typically, going out implies engaging in economic activity—work, study, business, etc. The individual payoffs from this choice depend on type, and on whether there is an immigration enforcement action r_{zt}, i.e., a raid, later in the day in the community. We summarize per-period payoffs with a utility function that mimics a 'tracking problem' (e.g., Baley and Veldkamp (2023)). For $i \in \tau$,

$$U^{\tau}_{izt}(a, r) = \pi^{\tau}(a, r) + \tilde{\xi}^{a}_{t} + \tilde{\zeta}^{a}_{zd(t)} + \epsilon^{a}_{izt}, \tag{10}$$

where $\tilde{\xi}^{a}_{t}$ and $\tilde{\zeta}^{a}_{zd(t)}$ are day and community-by-weekday common components, while the ϵ^{a}_{izt} are idiosyncratic unobservables, i.i.d. over i. The table below presents the type-specific component of payoffs:

$\pi^{N}(a, r)$	**Action**	
Raid	$a_{izt} = 0$	$a_{izt} = 1$
$r_{zt} = 0$	π_{SN}	π_{ON}
$r_{zt} = 1$	π_{SN}	π_{ON}

N-type payoffs

$\pi^{H}(a, r)$	**Action**	
Raid	$a_{izt} = 0$	$a_{izt} = 1$
$r_{zt} = 0$	π_{SN}	π_{ON}
$r_{zt} = 1$	π_{SR}	π_{OR}

H-type payoffs

Table of ex-post payoffs from the tracking game.

Because N-types are not subject to immigration enforcement, their payoff does not depend on whether a raid takes place or not, only on whether they stay home or go out. We will maintain $\pi_{ON} > \pi_{SN}$, so N-types only choose to stay home as a function of their idiosyncratic shock $(\epsilon^{0}_{izt}, \epsilon^{1}_{izt})$. The problem is more interesting for H-types. They would like

to coordinate their choice with the government's enforcement action. In the absence of a raid, they also rather go out (up to the idiosyncratic shock). If there is a raid, however, it is natural to assume they prefer to stay in, so $\pi_{SR} > \pi_{OR}$.

At the time of taking the action people do not know whether a raid will take place that day. However, they have a belief $B_{zd(t)} = \mathbb{P}(r_{zt} = 1|\mathcal{I}_{zt})$, common to everyone in the community, that a raid will happen. $\mathcal{I}_{zt}$ denotes the information available to the community at the beginning of the day. Note that we allow this belief to be weekday specific given our earlier discussion. This belief will only matter among H types since only they face immigration enforcement risk.

4.2.2 Pent-up demand

If staying home today increases the value of going out in the future, perhaps through a pent-up demand mechanism, the present reduction in activity may be fully or partially compensated by higher activity in subsequent days (see for example, Beraja and Zorzi (2024) for a similar mechanism in the context of durables consumption). If some of the contemporaneous decline in activity is driven by the perception of a likely raid today, this short-term chilling effect may be undone with above average activity in the future. To allow for this possibility we introduce a dynamic component of preferences for both H and N types. Each person accumulates 'pent-up demand for going out', x_{izt}. When $x > 0$, there is a backlog of unfulfilled desire to go out. We capture this idea modeling the evolution of pent-up demand as:

$$x_{izt+1} = (1 - \delta)x_{izt} - (a_{izt} - \kappa). \tag{11}$$

$\delta \in (0, 1)$ governs the rate at which pent-up demand decays over time, and $\kappa > 0$ is the desired long-run frequency of going out. Throughout we set $\kappa = 1$, so that staying in at time t generates an increase in pent-up demand of one day.[13]

Accumulated pent-up demand generates an instantaneous convex flow utility cost:

$$c(x_{izt}) = \frac{\psi}{2}x_{izt}^2, \qquad \psi \geq 0. \tag{12}$$

When $\psi = 0$ there is no cost of having refrained from going out in the past, and the going-out decision is static: low present consumption does not lead to higher future consumption. For $\psi > 0$, staying home today increases pent-up demand, making this a dynamic discrete choice

[13] $\kappa = 1$ has two implications: i) pent-up demand converges to $1/\delta$ for a person that never goes out. ii) for $x_{iz0} \geq 0$, pent-up demand cannot be negative, this is, people cannot find themselves 'ahead of schedule'.

problem. In this case, the comparison of payoffs when going out or staying home must take into account that staying home today, say because the raid belief is high, increases x_{izt+1}, which is costly next period.

4.2.3 Dynamic Programming Problem

We can express a type-τ person's dynamic problem at the beginning of period t with the following Bellman equation:

$$\begin{aligned} V^{\tau}(x_{izt}) =& \mathbb{E}_{\epsilon}\Bigg[\max_{a_{izt}\in\{0,1\}}\Bigg\{\mathbb{E}_{r}\Big[a_{izt}U^{\tau}_{izt}(1,r)+(1-a_{izt})U^{\tau}_{izt}(0,r)\Big|\mathcal{I}_{zt}\Big] \\ &-\frac{\psi}{2}x^{2}_{izt}+\beta V^{\tau}\big((1-\delta)x_{izt}+\kappa-a_{izt}\big)\Bigg\}\Bigg]. \end{aligned} \tag{13}$$

This is a standard dynamic discrete choice problem. Following Hotz and Miller (1993), assuming the idiosyncratic payoff components ϵ^{a}_{izt} are i.i.d type-I extreme value-distributed implies choice probabilities of the form

$$\mathbb{P}(a_{izt}=1|x_{izt},B_{zd(t)},\tau)=\frac{\exp(v^{\tau}_{1}(x_{izt},B_{zd(t)})-v^{\tau}_{0}(x_{izt},B_{zd(t)}))}{1+\exp(v^{\tau}_{1}(x_{izt},B_{zd(t)})-v^{\tau}_{0}(x_{izt},B_{zd(t)}))} \tag{14}$$

where we can write the value differences as

$$v^{\tau}_{1}(x_{izt},B_{zd(t)})-v^{\tau}_{0}(x_{izt},B_{zd(t)})=\alpha_{0}+\alpha_{1}B_{zd(t)}\mathbf{1}[\tau=H]+\xi_{t}+\zeta_{zd(t)}+\Delta^{\tau}(x_{izt})$$

for $\tau\in\{H,N\}$, and $\xi_{t}\equiv\tilde{\xi}^{1}_{t}-\tilde{\xi}^{0}_{t}$, $\zeta_{zd(t)}\equiv\tilde{\zeta}^{1}_{zd(t)}-\tilde{\zeta}^{0}_{zd(t)}$, $\alpha_{0}\equiv\pi_{ON}-\pi_{SN}>0$, $\alpha_{1}\equiv(\pi_{OR}-\pi_{SR})-(\pi_{ON}-\pi_{SN})$, and

$$\Delta^{\tau}(x_{izt})\equiv\beta\big[V^{\tau}((1-\delta)x_{izt}-(1-\kappa))-V^{\tau}((1-\delta)x_{izt}+\kappa)\big]. \tag{15}$$

Given the acute costs of being arrested for H types, absent any shocks it is natural to expect the net gain from going out relative to staying home if there is a raid, $\pi_{OR}-\pi_{SR}$, to be negative. And since the net gain from going out relative to staying home if there is no raid, $\pi_{ON}-\pi_{SN}$, is positive, $\alpha_{1}<0$. The component $\Delta^{\tau}(x_{izt})$ captures the dynamic utility gain from reduced pent-up demand into next period when going out instead of staying home.[14]

[14]The stationarity assumption shuts down anticipatory substitution, where agents also front-load consumption before expected high-risk days. This form of inter-temporal substitution would arise in a non-stationary version of our model where the value function depends on the entire future path of beliefs. By shutting down this channel, our estimates attribute any compensating dynamics entirely to pent-up demand. In particular, ψ would absorb some of the variation in the timing of economic activity, potentially overstating the pent-up demand cost and the associated welfare losses from enforcement. Relaxing stationarity would decompose the

4.2.4 Aggregation and Mixture Structure

In practice we work with data aggregated at the ZCTA level, so we estimate a version of the model above with a representative agent for each type, each of whom keeps track of the aggregate pent-up demand of his group, x^{τ}_{zt}. Thus, observed economic activity is

$$A_{zt} = h_z \Lambda^H_{zt}\left(w^H_{zt}\right) + (1 - h_z)\Lambda^N_{zt}\left(w^N_{zt}\right) \tag{16}$$

where $\Lambda(w)$ is the logistic function, and

$$w^{\tau}_{zt} = \alpha_0 + \alpha_1 B_{zd(t)}\mathbf{1}[\tau = H] + \Delta^{\tau}(x^{\tau}_{zt}) + s_{zt}, \quad s_{zt} \equiv \xi_t + \zeta_{zd(t)}.$$

The two latent stocks evolve as in (11) separately for H and N types. Notice that the date and weekday payoff shocks are common to both groups. The groups' aggregate choices differ through the impact of beliefs and the differential evolution of their aggregate pent-up demand. Equation (16) is a simple mixture logistic regression with observed weights, which depends on the following covariates: Beliefs, pent-up demand via the dynamic utility component, time fixed effects, and weekday fixed effects. We can estimate this regression solving for the value functions for each group as long as we can build a measure of beliefs.

4.3 Learning

4.3.1 Bayesian Learning about Enforcement Risk

We model how communities form and update beliefs about immigration enforcement risk. Our specification rests on two maintained assumptions. First, ICE's enforcement strategy does not depend on the community's decisions—neither economic activity nor belief updating affects when or where ICE conducts operations. Second, communities observe successful raids (those resulting in arrests) but not failed enforcement attempts. This assumption ensures that agents' information sets coincide with ours, which is necessary for constructing beliefs from observable data.

Given that raid outcomes are binary, we adopt a Beta-Bernoulli learning framework. Our reduced-form evidence motivates two refinements to standard Bayesian updating. First, because ICE enforcement exhibits pronounced day-of-the-week patterns, we allow agents to maintain weekday-specific beliefs: a Monday raid updates beliefs about future Mondays, not

compensating dynamics into backward-looking catch-up and forward-looking substitution, but at significant computational cost and with limited additional identification from the data, given the regularity of weekly enforcement patterns. Model validation, thus, is important to assess the empirical relevance of the pent-up demand mechanism.

about Tuesdays. Second, because raids exhibit positive serial correlation—a raid yesterday strongly predicts a raid today—we allow beliefs to condition on whether a raid occurred the previous calendar day. Together, these refinements mean agents effectively maintain fourteen parallel learning processes: one for each weekday × previous-day-raid-status combination.

Formally, let a_d^k and b_d^k denote the Beta parameters for weekday $d \in \{1, \ldots, 7\}$ conditional on previous-day raid status $r_{t-1} = k \in \{0, 1\}$. The prior mean for weekday d following $r_{t-1} = k$ is

$$B_{d(0)}^k = \frac{a_{d(0)}^k}{a_{d(0)}^k + b_{d(0)}^k}.$$

Standard Beta-Bernoulli updating implies that after observing the same-weekday outcome r_{t-7} from one week earlier, the parameters evolve as

$$\begin{bmatrix} a_{d(t)}^{r_{t-1}} \\ b_{d(t)}^{r_{t-1}} \end{bmatrix} = \begin{bmatrix} a_{d(t-7)}^{r_{t-1}} \\ b_{d(t-7)}^{r_{t-1}} \end{bmatrix} + \begin{bmatrix} \mathbf{1}\{r_{t-7} = 1\} \\ 1 - \mathbf{1}\{r_{t-7} = 1\} \end{bmatrix}. \tag{17}$$

The posterior mean belief about a raid on day t, conditional on weekday $d(t)$ and previous-day outcome r_{t-1}, is

$$B_{zd(t)} \equiv \mathbb{P}(r_t = 1 \mid r_{t-1}, r_{t-7}) = \frac{a_{d(t)}^{r_{t-1}} + \mathbf{1}[r_{t-7} = 1]}{a_{d(t)}^{r_{t-1}} + b_{d(t)}^{r_{t-1}} + 1}.$$

With learning incorporated, $B_{zd(t)}$ is no longer a free parameter but a deterministic function of the observed raid history $\{r_s\}_{s<t}$ and the initial prior.

4.3.2 Cumulative Forecast Error and Belief Sensitivity

The updating rule in (17) implicitly assumes individuals perceive ICE's strategy to be stationary. However, enforcement patterns may shift over time as ICE reallocates resources or adjusts priorities —just as we documented in Section 3.6. In such an environment, agents may prefer to weight recent signals more heavily than distant ones (see, e.g., Ma et al. (2023) in the computational cognitive science literature context, or the burgeoning literature in macroeconomics that studies expectation formation: Carvalho et al. (2023); Cho and Kasa (2015); Weber et al. (2025)).[15] We address this concern not by modifying the learning rule directly—which would introduce difficult-to-identify gain parameters—but through a belief-

[15] One could accommodate this with an updating rule of the form

$$\begin{bmatrix} a_{d(t)}^{r_{t-1}} \\ b_{d(t)}^{r_{t-1}} \end{bmatrix} = \lambda_t(\overline{f}_t) \begin{bmatrix} a_{d(t-7)}^{r_{t-1}} \\ b_{d(t-7)}^{r_{t-1}} \end{bmatrix} + \begin{bmatrix} \mathbf{1}\{r_{t-7} = 1\} \\ 1 - \mathbf{1}\{r_{t-7} = 1\} \end{bmatrix},$$

uncertainty interaction component described below, which attenuates behavioral responses when past beliefs have proven unreliable.[16]

As beliefs evolve through learning, agents accumulate a track record of forecast accuracy. We define the cumulative average forecast error as a weighted average of absolute forecast errors up to time t:

$$\overline{f}_{zt} = \frac{1}{t-1} \sum_{s=1}^{t-1} \omega_{st} |r_{zs} - B_{zd(s)}|, \tag{18}$$

where ω_{st} are exponential weights that down-weight older signals. $\overline{f}_{zt}$ captures how well agents' beliefs have tracked realized enforcement over the sample period. Notice that $\overline{f}_{zt}$ is *not* weekday-specific: we expect individuals to base their perceptions about regime changes on an overall assessment of forecast accuracy across all days, not on weekday-specific tallies.

A natural hypothesis is that the strength of the chilling effect varies with agents' confidence in their beliefs. When past forecasts have been accurate, agents should weight their beliefs heavily in decision-making. When forecasts have systematically missed the mark—perhaps because enforcement patterns shifted unexpectedly—agents may respond more cautiously to their current beliefs. To capture this, we allow the belief coefficient to be heterogeneous in forecast accuracy:

$$\alpha_1(\overline{f}_{zt}) = \alpha_{11} + \alpha_{12} \cdot \overline{f}_{zt}. \tag{19}$$

The H-type latent utility index thus becomes $w_{zt}^H = \alpha_0 + \left(\alpha_{11} + \alpha_{12} \cdot \overline{f}_{zt}\right) \cdot B_{zt} + c(x_{zt}^H) + s_{zt}$.

The parameter α_{11} captures the average effect of beliefs on the propensity to go out; we expect $\alpha_{11} < 0$ since higher perceived raid probability should reduce economic activity. The parameter α_{12} governs how this effect varies with forecast accuracy. If $\alpha_{12} > 0$, agents respond more weakly to their beliefs when past forecasts have been poor: the total belief effect $(\alpha_{11} + \alpha_{12} \cdot \overline{f}_{zt})$ becomes less negative as $\overline{f}_{zt}$ rises. Conversely, when forecasts have

where the gain parameter $\lambda_t < 1$ underweights past information as a function of previous forecast errors f_t. This is akin to the distinction between constant-gain and decreasing-gain learning in the macroeconomic expectations literature.

[16]The gain parameter is often weakly identified for several reasons: First, because it affects observables only indirectly, through the evolution of agents' beliefs, rather than entering the measurement equations directly. As a result, different combinations of learning speed and belief paths can generate very similar implications for observed outcomes, making he likelihood is often flat in the gain dimension, reflecting limited information in the data to pin down learning speed separately from other dynamic parameters. Second, the gain parameter is highly collinear with the persistence of structural shocks, since slow learning can mimic persistent exogenous processes. Third, beliefs are generated regressors inferred from the model rather than directly observed, which further attenuates identification of the gain parameter. This weak identification problem is well known in the macroeconomics literature on adaptive learning (Milani, 2007; Slobodyan and Wouters, 2012).

been accurate (low $\overline{f}_{zt}$), the chilling effect is stronger. This specification nests the baseline model: when $\alpha_{12} = 0$, the belief effect is constant at α_{11}.

This mechanism generates a sharp prediction: following structural breaks in enforcement patterns, we should observe spikes in $\overline{f}_{zt}$ as the break induces a mismatch between agents' learned beliefs and the new enforcement regime. We test this prediction directly in Section 6.

4.3.3 Discussion: Mis-specification of the Learning Model

A natural concern with any structural learning model is whether the assumed belief formation process is correctly specified.[17] We can identify several margins along which agents' actual learning might deviate from our specification. Our model addresses each of these concerns.

The first concern is non-stationarity: ICE's enforcement strategy may drift over time due to policy changes, resource reallocation, or political pressures, so that the underlying raid probability φ is not constant. Indeed we already documented the presence of distinguishable enforcement regimes over time. In such an environment, a standard Bayesian learner who weights all past signals equally would be slow to adapt to regime changes. Our proposal to allow for a heterogeneous effect of beliefs as a function of the cumulative forecast error directly addresses this concern. When accumulated forecast errors are large—signaling that the environment may have shifted—agents' response to their belief about enforcement risk is dampened. This is similar to the distinction between constant-gain and decreasing-gain learning in the macroeconomic expectations literature, but gets around the well known weak identification issues of such gain parameters at the cost of being somewhat reduced form.

The second concern is serial dependence: raid signals may be correlated over time in ways that a memoryless updating rule would miss. Our reduced-form evidence strongly supports this pattern—a raid yesterday substantially increases the probability of a raid today, independent of weekday. We incorporate this dependence by allowing agents to maintain separate tallies conditional on whether a raid occurred the previous day. Beliefs thus depend on $(d(t), r_{t-1})$, not just on the weekday alone.

Third, the underlying raid probability may be heterogeneous across contexts. The most salient heterogeneity in our setting is across weekdays: ICE enforcement exhibits pronounced day-of-the-week patterns. We accommodate this by specifying weekday-specific tallies.

Finally, the observed signal—whether a successful raid occurred—may be a noisy measure of the underlying enforcement activity. ICE may conduct operations that do not result in arrests and hence go unrecorded in our data. Crucially, agents in the community face the same informational constraint: they observe successful raids, not failed attempts. By

[17] See Manski (2004) for a discussion of these conceptual issues.

assuming that agents observe the same signal we do, we ensure consistency between the information set used to construct beliefs and the information agents actually have. This assumption rules out learning through informal networks about unsuccessful operations, but it is necessary for identification and conservative in the sense that any additional information agents possess would only sharpen their beliefs.

5 Estimation and Results

5.1 Sample and Estimation Strategy

Our estimation sample consists of the 352 ZCTAs in the top quintile of the Hispanic foreign-born share distribution that experienced at least 30 raids during the period (Oct. 2014 to May 2018). This threshold ensures sufficient within-ZCTA variation in beliefs and forecast errors. The resulting panel contains 471,328 ZCTA-day observations spanning 1,339 days.

We initialize beliefs using predicted raid probabilities from a probit regression estimated on a burn-in period of the 105 initial days of our sample. We only use the burn-in period observations for this purpose, and exclude them from the estimation sample, ensuring that beliefs and forecast errors have stabilized. The probit includes interactions of weekday and previous-day-raid status with Hispanic share and log median household income, as well as ICE Area of Responsibility fixed effects. Appendix F.2 provides more details and reports the probit estimates (Table A2).

We estimate the structural parameters $\boldsymbol{\theta} = (\psi, \delta, \alpha_{11}, \alpha_{12})$ by matching model-implied consumption to observed consumption, concentrating out the high-dimensional fixed effects. The approach exploits the invertibility of the mixture model: given $\boldsymbol{\theta}$, initial conditions for the prior B_{z0}, and for the pent-up demand state x_{z0}, and observables $(A_{zt}, B_{zt}, x_{zt}, h_z)$, we can recover a unique implied shock s_{zt} that rationalizes each observation. The estimation proceeds in three steps. First, for a candidate $\boldsymbol{\theta}$, we solve for the stationary value functions $V^\tau(x)$ and compute the dynamic incentive term $\Delta^\tau(x)$ from equation (15). Second, we invert the mixture model to recover the shocks $\{s_{zt}(\boldsymbol{\theta})\}$ that rationalize observed consumption.[18] Third, we regress these recovered shocks on date and ZCTA-by-weekday fixed effects. The estimator $\hat{\boldsymbol{\theta}}$ minimizes the residual sum of squares of this regression.

We fix the discount factor $\beta = 0.995$ (daily data) and the pent-up demand target rate $\kappa = 1$, and use an exponential function for the forecast error weights ω_s. The fixed effects absorb the intercept α_0, so we cannot separately identify it. Appendix F provides complete details on the solution algorithm, inversion procedure, and numerical implementation.

[18] This inversion yields a unique solution, as we describe in subsection F.3.

5.2 Identification

Dynamic parameters. The pent-up cost ψ and depreciation rate δ are identified from serial correlation in consumption after accounting for the fixed effects, following periods with changes in beliefs. When $\psi = 0$, the model collapses to myopic behavior: a belief shock ΔB_t affects only consumption on day t. When $\psi > 0$, pent-up states evolve smoothly, creating inter-temporal substitution (bounce-back). Thus, negative auto-correlation in economic activity —over calendar time— following a change in $B_{zd(t)}$ informs ψ. The positive auto-correlation in economic activity that follows while pent-up demand is still high then informs δ, since this parameter governs the horizon of its decay.[19]

Belief parameters. The belief effect α_{11} is identified from cross-sectional variation in the Hispanic foreign-born share h_z interacted with beliefs B_{zt}. High-h communities experience larger consumption drops one and seven days after a raid. The interaction parameter α_{12} is identified from time-varying sensitivity to beliefs within communities: periods of high vs. low forecast accuracy generate different consumption responses to the same belief level.

Notice that the belief effect and the pent-up effect work in *opposite directions* after a raid shock that shifts beliefs: while the increased risk perception causes consumption to drop (if $\alpha_{11} < 0$), accumulated pent-up demand causes consumption to rise later. This timing difference—immediate suppression followed by delayed recovery—allows separate identification of the two mechanisms.

5.3 Main Results

5.3.1 Main estimates

Table 3 presents the parameter estimates with three sets of standard errors: standard errors from the numerical Hessian, and Newey-West HAC standard errors with 8 and 30 lags.

Belief effect (α_{11}). The negative estimate $\hat{\alpha}_{11} = -0.28$ confirms that higher perceived raid risk reduces economic activity. A 10 percentage point increase in beliefs B_{zt} decreases the Hispanic foreign born 'going out' probability by 0.7 percentage points (evaluated at the sample mean).

[19] In contrast with previous literature on pent-up demand, where typically the depreciation parameter is set externally (e.g., Beraja and Zorzi (2024)), the variation available in our data together with the parsimony of the model allows us to estimate δ.

Belief-uncertainty interaction (α_{12}). The positive estimate $\hat{\alpha}_{12} = 0.93$ indicates that agents respond less strongly to their beliefs when past forecasts have been inaccurate. The total belief effect is $(\alpha_{11} + \alpha_{12} \cdot \bar{f}_{zt}) \cdot B_{zd(t)}$. At the mean cumulative forecast error ($\bar{f} \approx 0.07$), the effective belief coefficient is $-0.28 + 0.93 \times 0.07 = -0.21$, about 25 percent weaker than the baseline effect. At the 90th percentile of forecast error ($\bar{f} \approx 0.12$), the effective coefficient is $-0.28 + 0.93 \times 0.12 = -0.16$, nearly 40 percent weaker. This attenuation is intuitive: agents who have made poor predictions in the past place less weight on their current beliefs. The interaction mechanism provides a behavioral foundation for why chilling effects may be heterogeneous across time and communities.

Pent-up demand (ψ, δ). The estimates $\hat{\psi} = 2.233$ and $\hat{\delta} = 0.417$ imply meaningful inter-temporal substitution. When chilling effects suppress consumption, pent-up demand accumulates and eventually manifests as above-baseline activity. The depreciation rate implies a half-life of $\log(2)/\log(1/(1-0.417)) \approx 1.3$ days, suggesting that pent-up effects dissipate relatively quickly. Figure A13 in Appendix F.1 displays the estimated value function and the dynamic incentive term $\Delta(x)$, illustrating how the option value of going out varies with the pent-up state.

Fixed effects (ξ, ζ). The estimation concentrates out 1,234 date fixed effects and 2,464 ZCTA-by-weekday fixed effects. Figure A16 in Appendix F displays these estimated fixed effects: the date effects capture aggregate time-series variation (holidays, macroeconomic shocks) and reveal a downward sloping trend in aggregate activity over time in the sample, while the ZCTA-by-weekday effects illustrate meaningful level differences in aggregate activity between ZCTAs across weekdays.

Implied beliefs and forecast errors. Figure A17 displays the implied beliefs, cumulative forecast errors, and time-varying marginal effect of beliefs. Beliefs exhibit the expected weekday structure, with higher raid probabilities on weekdays than weekends. The distribution of marginal effects of beliefs, $\alpha_{11} + \alpha_{12} \cdot \bar{f}_{zt}$, shows mean shifts at different points in time but is overall stationary. The most pronounced mean shifts take place around the times when there is a large number of ZCTAs for which we detected regime changes. The time variation in this marginal effect captures, through the lens of our model, a similar mechanism to the decreasing gain term in models of adaptive learning, and is a quantitatively important source of variation in economic activity. Figure A18 in Appendix F.8 additionally illustrates implied mean raid probabilities and the evolution of beliefs for a handful of ZCTAs.

5.3.2 Impulse Response Analysis

To investigate further the model's dynamic implications, we now discuss a series of impulse-response exercises based on our estimates. Starting from steady state, we trace how economic activity responds to an unexpected increase in beliefs—the dynamic analog of the chilling effect. We do so simulating 'going out' choices forward from steady state under two scenarios: i) a shock to beliefs followed by immediate reversion back to baseline, and ii) a shock to beliefs after which beliefs are allowed to evolve forward in response to the shock. The impulse response measures the percent difference in economic activity relative to steady state. In Figure 7 we present the results of these exercises under a 15 percentage point shock to beliefs at $t = 0$. This is a large shock, as estimated mean beliefs outside weekends are between 6 to 10 percent. We report results on mean overall economic activity, and on mean economic activity among the Hispanic foreign born.

On impact, the belief shock causes economic activity to fall, with a much larger effect for the Hispanic foreign-born group (approximately 1 percent of their economic activity) than for the aggregate (approximately 0.14 percent of overall economic activity). This differential reflects that only the H-type group responds to enforcement beliefs and that the Hispanic foreign born are 14 percent of the typical ZCTA population in our estimation sample.

Following the initial drop, pent-up demand dynamics generate a rapid rebound. The top-left panels of Figure 7 show the pent-up state x rising on impact as suppressed activity accumulates, then gradually declining as agents catch up. By day 1, economic activity is already above baseline as the suppressed activity from day 0 spills forward. The pent-up demand half-life of approximately 1.3 days (implied by $\hat{\delta} = 0.42$) means most of the rebound occurs within the first three days.

Despite the rebound, the cumulative effect remains negative: the model predicts that pent-up demand does not fully offset the initial suppression. Only around half of the impact economic activity loss is compensated for in later days. This finding is consistent across both scenarios—whether beliefs revert immediately or evolve endogenously after the shock.

Figure A19 then examines impulse responses across shock magnitudes between 5 and 70 percentage points. The impact and cumulative responses are approximately linear. Notably, the scenario where beliefs evolve endogenously (red lines) shows slightly larger cumulative losses than the immediate-reversion scenario (blue lines), reflecting the persistence of belief shocks when updating is allowed. The structural impulse responses are qualitatively consistent with our reduced-form event studies, which also show an immediate consumption drop followed by recovery. The structural model, however, allows us to distinguish between the short-run and cumulative chilling effect, highlighting that impact effects may overstate the welfare losses when individuals can shift their consumption over time.

6 Model Validation

In this section we assess the model's performance through goodness-of-fit diagnostics and—most importantly—an exercise using the structural breaks for model validation.

6.1 Model Fit

Figure A20 presents four diagnostics that assess the model's fit across both cross-sectional and time-series dimensions. Panel (a) shows a binned scatter plot of predicted versus observed economic activity. The model explains approximately 12 percent of the variation in consumption after removing fixed effects ($R^2 = 0.118$, RMSE = 0.042). Panel (b) shows a binned scatter on the cross-section of ZCTAs, plotting implied mean belief $B_{zd(t)}$ against the observed raid frequencies. Beliefs slightly under-predict raid probabilities at high levels, but the overall correspondence is strong. Panel (c) then shows how model fit varies with the Hispanic foreign-born share. The RMSE is relatively stable across the distribution, indicating that the model does not systematically fit better or worse in communities with different demographic compositions. This is reassuring given that part of the model parameter's identification relies on cross-sectional variation in the Hispanic share. Finally, panel (d) displays the distribution of model residuals over time. They are well-centered around zero throughout the sample period, albeit with a slight downward drift over time. The spikes in the plot, which happen around every 30 days, reveal beginning-of-the-month seasonality in economic activity, which our model did not control for. Within ZCTA, residuals do exhibit some autocorrelation, particularly at short lags as we show in Appendix F.6, Figure A15. This motivates our reporting of HAC standard errors for the model parameter estimates.

6.2 Structural Breaks as Model Validation

Our model implies a specific prediction: when ICE enforcement patterns change unexpectedly, agents' beliefs will temporarily mis-predict raids, causing forecast errors to spike. We test this prediction using the structural breaks identified in Section 3.6 using the Bai and Perron (2003) methodology for the 352 ZCTAs in our estimation sample. We detect 233 structural breaks across 182 ZCTAs (52 percent have at least one break). The median regime duration is 503 days, indicating that enforcement patterns are reasonably stable but do shift periodically.

Figure 8 presents two pieces of evidence that validate our model's learning dynamics. Panel (a) displays an event study of mean forecast errors around detected structural breaks. Forecast errors are stable in the 120 days before the break, then spike sharply at event time

zero. The increase is substantial—approximately 73 percent above the pre-break level—and highly statistically significant. The spike then gradually decays as agents' beliefs adjust to the new regime, returning to near-baseline levels after approximately a year. This pattern is what our model predicts: when the world changes in ways agents have trouble anticipating, their prediction-accuracy deteriorates. Because our structural model does not know anything about the structural breaks—which are themselves estimated separately—, we find this validation test remarkable.[20]

Panel (b) shows the distribution of regime changes by direction. We find that 149 of 182, or 82 percent of detected breaks, represent *escalations* in enforcement intensity—an increase in the underlying raid probabilities—rather than de-escalations. The mean increase is 6.5 percentage points in the daily raid probability. The asymmetry is notable: while de-escalations are larger in magnitude when they occur (mean decrease of 13.7 percentage points), they are far less common. This pattern is consistent with a general trend toward intensified enforcement during our sample period.

6.3 Social Cohesion Underlies Learning

Our model estimates a community-level learning process but does not take a stand on the micro-level mechanisms behind it. We expect information diffusion within the local social networks to play a first order role: residents share experiences, warn neighbors, etc. If this is the case, communities with tighter, more interconnected networks should learn more effectively. Our model, which uses no information on such networks, should better capture the relationship between risk perception and economic activity in those ZCTAs.

We test this by regressing the average RMSE of the model residuals at the ZCTA level on three Facebook-based measures of social cohesion: the clustering coefficient and support ratio from Chetty et al. (2022), which measure how densely interconnected residents' friendship networks are, and the within-ZCTA Social Connectedness Index from Bailey et al. (2018). Table A3 reports the results, controlling for log population and log number of raids. All coefficients are negative, with the clustering coefficient and the support ratio statistically significant. The first principal component of the three measures is also significant. These patterns are consistent with local social networks facilitating the community-level learning we document.

[20] In Appendix G, Figure A21 we present the results of a placebo exercise where we randomly reassign break dates. Placebo breaks show no effect on forecast errors, confirming that the belief dynamics that our model recovers reflect genuine changes in the underlying enforcement patterns.

7 Counterfactuals

The structural model provides a natural framework for counterfactual analysis. We now present the results of a series of simulations of economic activity under alternative enforcement scenarios. These simulations allow us to ask questions about the value of information and learning, and to assess the welfare consequences of different enforcement policies.

7.1 Methodology

Given our estimated parameters $\hat{\theta} = (\hat{\psi}, \hat{\delta}, \hat{\alpha}_{11}, \hat{\alpha}_{12})$ and the structural shocks $\hat{s}_{zt}$ recovered from the data, we simulate economic activity under alternative enforcement scenarios by solving the model forward. The key feature of this approach is that it captures two sources of endogenous dynamics simultaneously. Pent-up demand evolves as suppressed activity today accumulates into future consumption pressure, and beliefs evolve as agents update their posteriors in response to the enforcement signals they observe.

These dynamics are central to interpreting our results, and have first-order consequences for any welfare assessment of enforcement policies that generate chilling effects. This realistic treatment of transition dynamics distinguishes our counterfactuals from comparative statics exercises that assume instantaneous adjustment. For scenarios involving stochastic elements—such as randomly removing a fraction of observed raids—we use Monte Carlo simulations and report means over the draws.

7.2 Willingness to Pay and Immigration Enforcement

7.2.1 The Value of Perfect Information about Enforcement Activity

We begin with an exercise to quantify the value of information. To answer this question we use our model estimates to simulate economic activity in a scenario of perfect foresight about immigration enforcement. Thus, we suppose agents know whether a raid will occur before making their going-out decision, so $B_{zd(t)} = r_{zt}$. This eliminates forecast errors: $\overline{f}_{zt} = 0$ for all z, t. By measuring by how much economic activity changes compared to economic activity under the observed enforcement risk, we obtain an implicit measure of how much economic activity people in these communities would be willing to give up to be perfectly informed about enforcement actions.

Figure 9 presents the results. Panel (a) shows the distribution of consumption changes for H-types across ZCTAs. The mean effect is -1.7 percent: under perfect information, economic activity *falls* by about 0.6 days per year (in the data, the average yearly 'going-out' rate is 35 days). This decrease reflects the value agents place on being able to optimally

time their behavior. With perfect foresight, agents stay home exactly when raids occur and go out exactly when they do not—they avoid enforcement risk at the correct rate. The consumption they forgo to achieve this optimal timing can be interpreted as what they would be willing to pay for perfect information.

Panel (b) shows that this willingness to pay varies systematically with forecast accuracy under the baseline. ZCTAs with higher cumulative forecast errors—those whose beliefs were less accurate—would pay more for perfect information. Communities in the lowest forecast error quartile would sacrifice only 0.2 percent of consumption, while those in the highest quartile would sacrifice 3.0 percent. Thus, communities who have been forecasting poorly have the most to gain from knowing enforcement outcomes in advance.

The finding that 89 percent of ZCTAs show consumption decreases indicates that most communities would willingly pay something for perfect information. The average willingness to pay of 1.7 percent of H-type consumption—modest but meaningful—reflects the welfare cost of living under enforcement uncertainty.

7.2.2 The Value of No Immigration Enforcement

We now ask a different question: how much would people be willing to pay to live in a world with no immigration enforcement at all? To answer this, we compare economic activity under two scenarios. In the first, we set $B_{zd(t)} = 0$ throughout—agents believe no raids will ever occur and go out at their unconstrained optimum. In the second, we set $B_{zd(t)} = \bar{r}_{z,d}$, the average raid probability for ZCTA z on weekday d—agents face the true average raid risk with certainty. Crucially, in both scenarios there is no idiosyncratic risk. This ensures an apples-to-apples comparison: we isolate the welfare cost of facing raid risk itself, holding constant the accuracy of beliefs.[21]

The consumption difference between these scenarios measures how much economic activity agents would sacrifice to eliminate enforcement risk entirely. Under no enforcement ($B = 0$), consumption is higher because agents face no deterrence. Under known risk ($B = \bar{r}$), consumption is lower because agents optimally reduce activity to avoid enforcement exposure. The gap represents the willingness to pay—in consumption terms—for a world without immigration enforcement.

Figure 10 presents the results. Panel (a) shows the distribution of consumption sacrifice among the Hispanic foreign born across the 352 ZCTAs. The mean sacrifice is 3.6 percent of their average consumption (1.2 'staying home' days out of the average 33 'going out' days

[21] Naturally, an environment without immigration enforcement would likely see changes in patterns of immigration, residential sorting, labor supply, etc. The implications of the exercise here are only partial equilibrium.

per year). This represents a substantial and persistent cost: even when agents know the true raid probability with certainty, the mere existence of enforcement risk causes them to reduce economic activity by nearly 4 percent relative to the world where they face none.

Naturally, observed enforcement intensity drives significant heterogeneity in the value of living under no immigration enforcement risk: the scatter plot in panel (b) illustrates that consumption sacrifice rises steeply with the observed average raid probability (correlation coefficient of 0.66). In contrast, the correlation with the Hispanic foreign-born share is essentially zero (-0.04). This pattern indicates that enforcement intensity, not demographic composition, determines the welfare costs of living under immigration enforcement.

The 3.6 percent sacrifice we estimate provides a lower bound on the total welfare cost of enforcement, as it captures only the anticipatory behavioral response on the consumption margin, and not the direct welfare costs from being arrested, or the psychological costs from living in an environment of persistent uncertainty.

In these exercises, agents *know from the start* that enforcement risk is either zero or at its average level; there is no uncertainty about the regime. Thus, there is no learning. If instead we implement an exercise where agents begin with beliefs calibrated to historical enforcement and must learn over time that raids have stopped, we find yearly consumption gains of 1.5 percent relative to the same known-aggregate risk scenario. This is around half the 3.6 percent we estimated under the known no-immigration enforcement world. The ratio of realized gains to potential gains—approximately 42 percent (1.5/3.6)—quantifies how much of the welfare benefit from eliminating enforcement can be captured during a realistic transition period of about a year. The remaining 58 percent represents the cost of learning frictions: even after enforcement stops, it takes time for beliefs to fully adjust.

7.3 Raid Likelihood Increase Scenarios

We next consider what happens if enforcement intensity increases permanently. We simulate scenarios where raid probabilities are scaled up by 25, 50, 75, and 100 percent relative to baseline levels, maintaining each ZCTA's weekday-specific enforcement patterns. Beliefs are initialized at steady-state values matching the scaled probabilities, allowing us to isolate the direct effect of higher enforcement risk on economic activity.

Table 4 summarizes the consumption effects for the Hispanic foreign born across all ZCTAs. The losses are monotonically increasing in the increase level: a 25 percent increase in raid probability yields a 1.0 percent reduction in economic activity, rising to 2.0 percent at 50 percent increase, 3.0 percent at 75 percent increase, and 4.0 percent when raid probabilities are doubled. Appendix Figure A22 reports the full distribution of losses across ZCTAs,

revealing these are bimodal.

7.4 The Consequences of Discounting Noisy Signals

A key feature of our model is that agents discount the chilling effect of beliefs when those beliefs have been inaccurate. This is captured by the interaction term $a_2 > 0$: when cumulative forecast errors are large, the marginal effect of beliefs on behavior is attenuated. To quantify the implications of this behavioral response, we compare economic activity under our estimated model ($\hat{a}_2 = 0.93$) to a counterfactual where agents are standard Bayesian learners who respond to beliefs with constant intensity ($a_2 = 0$).

Figure 11 presents the results. Panel (a) plots the percent increase in economic activity under our baseline estimates relative to economic activity in the "standard learning" scenario, against the Hispanic foreign-born share. The correlation is essentially zero (0.09), indicating that demographic composition plays almost no role in determining the consumption impacts of discounted learning. Panel (b) shows the distribution of these consumption gains. The mean gain among the Hispanic foreign born is 1.85 percent (0.72 days per year), with gains ranging from near zero to over 5 percent. This represents a substantial difference in activity driven by the discounting of unreliable signals.

The variation across ZCTAs is almost entirely driven by forecast error exposure, as expected. Panel (a) illustrates that ZCTAs facing very predictable enforcement patterns (bluest dots in the panel) gain little from belief discounting, while those facing hard to predict patterns (yellowest dots in the panel) gain considerably. Notice that the behavioral response we modeled here makes agents who discount inaccurate beliefs have higher consumption than if they did not discount them. This may or may not improve their welfare since whether higher consumption is or not desirable depends on whether current beliefs are below or above the true underlying raiding likelihood, something agents do not know.

8 Conclusions

This paper studies how immigrant communities learn to anticipate and adapt to immigration enforcement. We combine daily bank account transactions data with arrest-level records of ICE operations, and document that enforcement follows predictable weekday patterns and that communities can learn them. Consumption tracks expected raid likelihood, not just realized raids, showing that behavior responds to beliefs about enforcement risk.

We propose a structural model of pent-up demand with Bayesian learning that allows us to decompose chilling effects into consumption that is merely shifted to safer days versus

consumption that is permanently foregone. Roughly half of the immediate suppression is recovered through subsequent bounce-back; the other half represents genuine welfare loss. Counterfactual exercises reveal that eliminating enforcement risk would be worth 3.6 percent of Hispanic foreign-born consumption. However, when enforcement is removed, only 42 percent of these potential gains materialize during a realistic transition—the remainder reflects learning frictions as beliefs adjust to the new environment.

These findings have several implications. For a welfare perspective, they suggest that the benefits of reduced enforcement for households affected by deportation risk depend not only on the magnitude of the reduction but also on how quickly affected populations can learn that their environment has changed. Gradual reductions may generate smaller welfare gains than their steady-state effects would suggest. For the literature on enforcement and "chilling effects," our results highlight that short-run estimates using realized enforcement as a proxy for beliefs suffer substantial attenuation bias, but that inter-temporal consumption shifting from high to low enforcement risk days makes the unattenuated responses, which we estimate using an IV strategy, over-estimate the long-run effects.

More broadly, our findings document a form of informational adaptation to state power. Populations living under enforcement risk develop forecasting capacity, learning to read patterns in state action and adjusting their behavior accordingly. This adaptation partially mitigates the costs of enforcement uncertainty—but the costs remain substantial, and the learning process itself imposes welfare losses during policy transitions.

Our approach can be applied in settings where individuals repeatedly face uncertainty and where event data exhibit predictable structure. Examples include how households adjust spending or mobility in response to crime risk, how criminals respond to police patrolling intensity, how firms respond to predictable patterns of inspections or regulatory enforcement, or how individuals adjust mobility, human capital investments (schooling, healthcare utilization) and market participation in response to recurring conflict or violence.

References

Abadie, A. and J. Gardeazabal (2003): "The economic costs of conflict: A case study of the Basque Country," *American Economic Review*, 93, 113–132.

Acemoglu, D. and J. A. Robinson (2019): *The Narrow Corridor: States, Societies, and the Fate of Liberty*, New York: Penguin Press.

Alfano, M. and J.-S. Görlach (2022): "Terrorism, Media Coverage, and Education:

Evidence from al-Shabaab Attacks in Kenya," *Journal of the European Economic Association*, 21, 727–763.

Alsan, M. and C. S. Yang (2024): "Fear and the safety net: Evidence from secure communities," *Review of Economics and Statistics*, 106, 1427–1441.

Amuedo-Dorantes, C. and F. M. Antman (2022): "De facto immigration enforcement, ICE raid awareness, and worker engagement," *Economic Inquiry*, 60, 373–391.

Bai, J. and P. Perron (2003): "Computation and Analysis of Multiple Structural Change Models," *Journal of Applied Econometrics*, 18, 1–22.

Bailey, M., R. Cao, T. Kuchler, J. Stroebel, and A. Wong (2018): "Social connectedness: Measurement, determinants, and effects," *Journal of Economic Perspectives*, 32, 259–280.

Baley, I. and L. Veldkamp (2023): "Bayesian learning," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw, Academic Press, 717–748.

Beraja, M. and N. Zorzi (2024): "Durables and Size-Dependence in the Marginal Propensity to Spend," NBER Working Paper 32080, National Bureau of Economic Research.

Besley, T. and H. Mueller (2012): "Estimating the peace dividend: The impact of violence on house prices in Northern Ireland," *American Economic Review*, 102, 810–833.

Byrd, R. H., P. Lu, J. Nocedal, and C. Zhu (1995): "A Limited Memory Algorithm for Bound Constrained Optimization," *SIAM Journal on Scientific Computing*, 16, 1190–1208.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023): "Anchored inflation expectations," *American Economic Journal: Macroeconomics*, 15, 1–47.

Chalfin, A. and J. McCrary (2017): "Criminal deterrence: A review of the literature," *Journal of Economic Literature*, 55, 5–48.

Chetty, R., M. O. Jackson, T. Kuchler, J. Stroebel, N. Hendren, R. B. Fluegge, S. Gong, F. Gonzalez, A. Grondin, M. Jacob, D. Johnston, M. Koenen, E. Laguna-Muggenburg, F. Mudekereza, T. Rutter, N. Thor,

W. Townsend, R. Zhang, M. Bailey, P. Barberá, M. Bhole, and N. Wernerfelt (2022): "Social capital I: measurement and associations with economic mobility," *Nature*, 608, 108–121.

Cho, I.-K. and K. Kasa (2015): "Learning and Model validation," *The Review of Economic Studies*, 81, 45–82.

Ciancio, A. and C. García-Jimeno (2024): "The political economy of immigration enforcement: Conflict and cooperation under federalism," *Review of Economics and Statistics*, 106, 1460–1476.

Conley, T. G. and C. R. Udry (2010): "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, 100, 35–69.

Draca, M., S. Machin, and R. Witt (2011): "Panic on the streets of London: Police, crime, and the July 2005 terror attacks," *American Economic Review*, 101, 2157–2181.

East, C. N., A. L. Hines, P. Luck, H. Mansour, and A. Velásquez (2023): "The labor market effects of immigration enforcement," *Journal of Labor Economics*, 41, 957–996.

Evans, G. W. and S. Honkapohja (2001): *Learning and expectations in macroeconomics*, Princeton University Press.

García-Jimeno, C. (2016): "The political economy of moral conflict: An empirical study of learning and law enforcement under prohibition," *Econometrica*, 84, 511–570.

Gardner, J., N. Thakral, L. T. Tô, and L. Yap (2024): "Two-Stage Differences in Differences," Working Paper, Boston University.

Gonçalves, F. M., E. Jácome, and E. K. Weisburst (2024): "Immigration enforcement and public safety," Tech. rep., National Bureau of Economic Research.

Hotz, V. J. and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 60, 497–529.

Lochner, L. (2007): "Individual Perceptions of the Criminal Justice System," *American Economic Review*, 97, 444–460.

Ma, W. J., K. P. Kording, and D. Goldreich (2023): *Bayesian Models of Perception and Action: An Introduction*, Cambridge, MA: The MIT Press.

Malmendier, U. and S. Nagel (2016): "Learning from inflation experiences," *The Quarterly Journal of Economics*, 131, 53–87.

Manski, C. F. (2004): "Measuring Expectations," *Econometrica*, 72, 1329–1376.

Milani, F. (2007): "Expectations, learning and macroeconomic persistence," *Journal of monetary Economics*, 54, 2065–2082.

Munshi, K. (2004): "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution," *Journal of Development Economics*, 73, 185–213.

Newey, W. K. and K. D. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Orphanides, A. and J. C. Williams (2005): "Inflation scares and forecast-based monetary policy," *Review of Economic Dynamics*, 8, 498–527.

Pastor, L. and P. Veronesi (2012): "Uncertainty about government policy and stock prices," *The Journal of Finance*, 67, 1219–1264.

Pew Research Center (2025): "U.S. Unauthorized Immigrant Population Reached a Record 14 Million in 2023," .

Rubin, D. B. (1981): "The Bayesian Bootstrap," *The Annals of Statistics*, 9, 130–134.

Scott, J. C. (1985): *Weapons of the Weak: Everyday Forms of Peasant Resistance*, New Haven: Yale University Press.

——— (2009): *The Art of Not Being Governed: An Anarchist History of Upland Southeast Asia*, New Haven: Yale University Press.

Slobodyan, S. and R. Wouters (2012): "Learning in a medium-scale DSGE model with expectations based on small forecasting models," *American Economic Journal: Macroeconomics*, 4, 65–101.

Tapsoba, A. (2023): "The cost of fear: Impact of violence risk on child health during conflict," *Journal of Development Economics*, 160, 102975.

Watson, T. (2014): "Inside the refrigerator: Immigration enforcement and chilling effects in Medicaid participation," *American Economic Journal: Economic Policy*, 6, 313–338.

Watson, T. and K. Thompson (2022): "The border within: The economics of immigration in an age of fear," in *The Border Within*, University of Chicago Press.

Weber, M., B. Candia, H. Afrouzi, T. Ropele, R. Lluberas, S. Frache, B. Meyer, S. Kumar, Y. Gorodnichenko, D. Georgarakos, O. Coibion, G. Kenny, and J. Ponce (2025): "Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings," *Econometrica*, 93, 229–264.

Tables

	Below Median HF		Above Median HF		
	Mean	Std. Dev.	Mean	Std. Dev.	p-value
Panel A: Economic Activity					
Share accounts active	0.105	0.031	0.106	0.028	0.206
Spent amount ($)	1,756	4,306	2,526	5,375	0.000
Number of accounts	129.0	112.7	183.4	140.4	0.000
Accounts per 1,000 pop.	8.263	3.500	7.038	3.383	0.000
Panel B: Demographics					
Hispanic foreign-born share	0.006	0.005	0.078	0.079	0.000
Median household income ($)	65,787	28,124	65,897	28,164	0.824
Population	16,394	12,138	27,665	17,920	0.000
Panel C: ICE Enforcement					
Any raid (daily)	0.067	0.148	0.185	0.242	0.000
Raid count (cond. on raid)	1.401	0.655	1.873	1.097	0.000
Ever raided	0.635	—	0.841	—	0.000
N (ZCTAs)	6,433		6,432		
N (Days)	1,339				

Table 1: Summary Statistics

Notes: Summary statistics by median Hispanic foreign-born (HF) share. Panel A reports economic activity from Facteus. Panel B reports ZCTA characteristics from 2018 ACS 5-Year estimates. Panel C reports ICE enforcement from TRAC. Last column reports p-values for difference of means tests. Sample period: October 2014 – May 2018 (1,339 days).

	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV Overid
Panel A: Second Stage					
Raid Indicator	-0.000	-0.007	-0.012	-0.036	-0.009
	(0.000)	(0.003)	(0.002)	(0.005)	(0.003)
Panel B: First Stage					
Instrument (13-week avg)		0.286	0.425	0.592	
		(0.009)	(0.007)	(0.006)	
Raid Lag 1 Week					0.051
					(0.002)
Raid Lag 2 Weeks					0.044
					(0.002)
ZCTA × Weekday FE	Yes	Yes			Yes
ZCTA FE			Yes		
Weekday FE			Yes		
Date FE	Yes	Yes	Yes	Yes	Yes
Observations	1,490,112	1,490,112	1,490,112	1,490,112	1,582,050
R^2	0.56	0.56	0.52	0.18	0.56
First-stage F-stat		13,377	36,078	95,517	3,784

Table 2: Effect of Raid Beliefs on Economic Activity

Notes: The dependent variable is the share of accounts active in a ZCTA-day. The raid indicator equals one if at least one ICE arrest occurred in the ZCTA on that day. Standard errors clustered by ZCTA × weekday in parentheses. Columns 2–4 instrument for the raid indicator using the 13-week rolling mean of lagged same-weekday raids. Column 5 uses 1-week and 2-week same-weekday lags as separate instruments. Sample restricted to ZCTAs in the top quintile of the Hispanic foreign-born share distribution.

Parameter	Estimate	SE (Std)	SE (HAC-8)	SE (HAC-30)
ψ	2.23	0.01	0.09	0.09
δ	0.42	0.00	0.01	0.01
α_{11}	-0.28	0.02	0.12	0.12
α_{12}	0.93	0.04	0.52	0.51

Table 3: Pent-up Demand Model Parameter Estimates

Notes: 434,368 observations across 352 ZCTAs. SE (Std) reports standard errors from the numerical Hessian. SE (HAC-8) and SE (HAC-30) report Newey-West heteroskedasticity and autocorrelation-consistent standard errors with 8 and 30 lags, respectively.

Scenario	Mean Loss (days/year)	Mean Loss (%)	Median Loss (%)
Baseline (0%)	—	—	—
25% increase	0.40	1.02	1.09
50% increase	0.79	2.02	2.16
75% increase	1.16	3.00	3.22
100% increase	1.53	3.95	4.24

Table 4: Consumption Losses Under Increased Enforcement

Notes: Consumption losses among the Hispanic foreign born under counterfactual scenarios where raid probabilities are scaled up by k percent. Beliefs are initialized at steady-state values matching the scaled probabilities. Losses are computed relative to the baseline (k=0) and averaged over 50 Monte Carlo simulation draws.

Figures

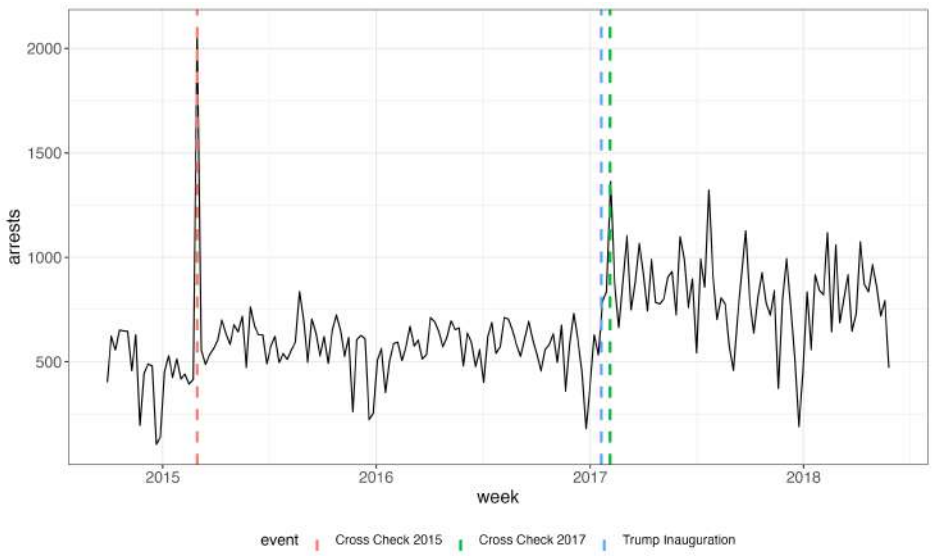


Figure 1: Weekly ICE Interior Arrests, October 2014 – May 2018.

Notes: Weekly count of ICE interior raid arrests from TRAC data. Red and green vertical lines indicate Operations Crosscheck (March 2015 and January 2017) and President Trump's inauguration (January 2017). Interior arrests exclude arrests by CBP and transfers from other agencies.

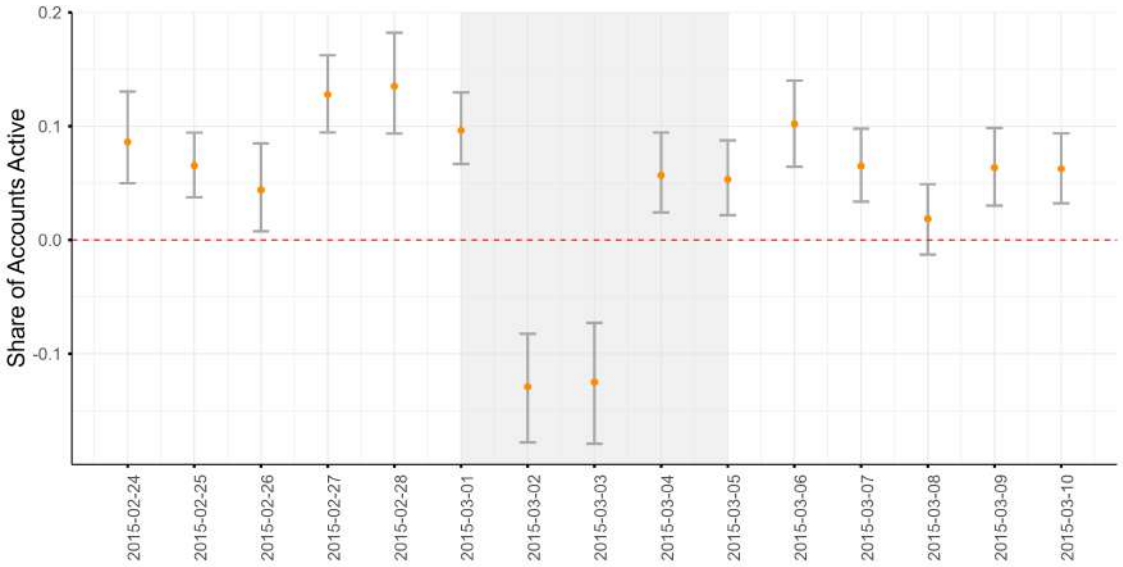
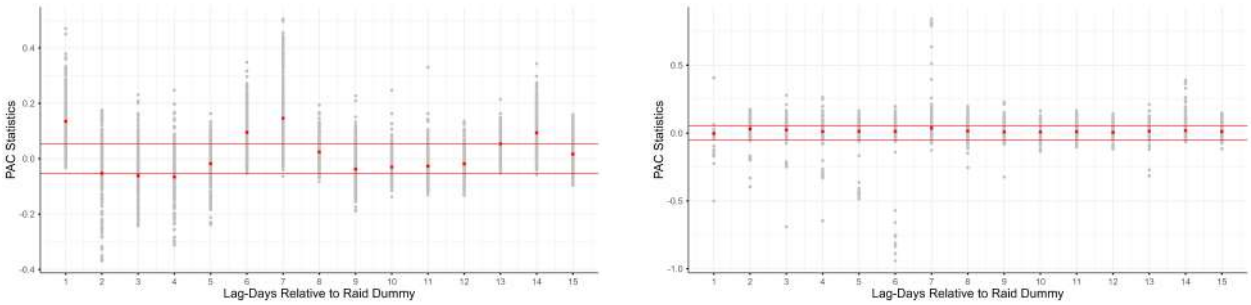


Figure 2: Operation Crosscheck Event Study: Hispanic Foreign-Born Share Interactions

Notes: Event study coefficients on the interaction between event time indicators and Hispanic foreign-born share. The sample includes 802 ZCTAs that experienced a raid during Operation Crosscheck (March 1–5, 2015) with no raids in the 10 days before or after. Event time is relative to March 3, 2015 (Tuesday of Crosscheck week). The estimation uses the Gardner et al. (2024) two-stage method with ZCTA × weekday and date fixed effects. Error bars show 95 percent confidence intervals from 1,000 Bayesian bootstrap replications (Rubin, 1981) clustered at the ZCTA level.



(a) All observations **(b)** Residualized

Figure 3: Partial Autocorrelation of ICE Raids

Notes: Panel (a) shows the partial autocorrelation function (PACF) of the raid indicator. Panel (b) shows the PACF of the residualized raid indicator, after controlling for weekday dummies and a lagged raid indicator. Red points indicate mean PACF values across ZCTAs; gray points show individual ZCTA values.

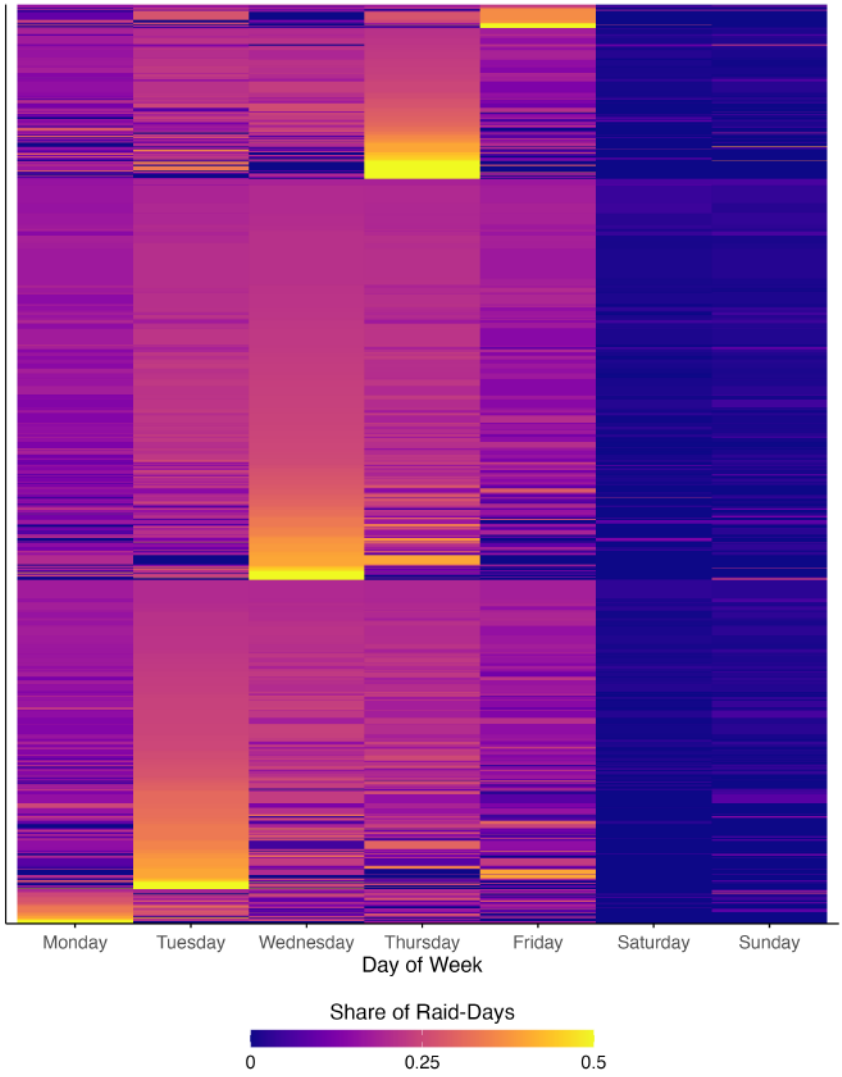
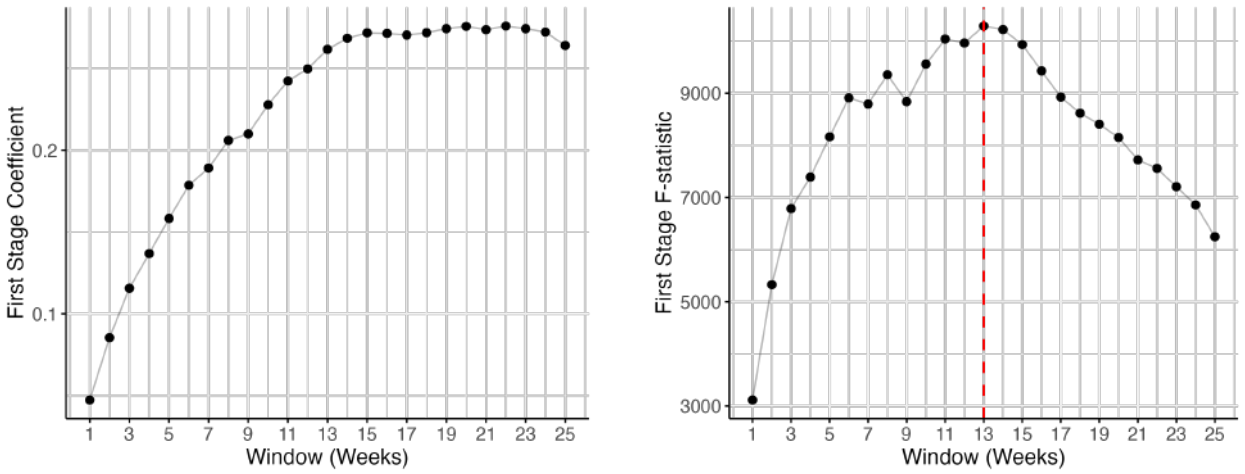


Figure 4: Heterogeneity in Weekday Enforcement Patterns Across ZCTAs

Notes: Heatmap of weekday enforcement profiles for 6,711 ZCTAs with raids between 2014 and 2018. Each row represents one ZCTA. Each column shows the share of that ZCTA's raid-days occurring on each weekday. ZCTAs are grouped by peak enforcement day and sorted by concentration within each group. While mid-week enforcement is generally higher (vertical structure), the intensity varies substantially across ZCTAs (horizontal variation within each block).



(a) First Stage Coefficient

(b) First Stage F-statistic

Figure 5: Instrument Selection: First-Stage Statistics by Window Length

Notes: First-stage statistics from IV regressions using the rolling mean of lagged same-weekday raids as instrument for the current raid indicator. Panel (a) shows first-stage coefficient; panel (b) shows first-stage F-statistic. The vertical dashed line in panel (b) indicates the optimal window (13 weeks) by F-statistic. All regressions include ZCTA × weekday and date fixed effects. Standard errors clustered by ZCTA × weekday. The sample is restricted to ZCTAs in top quintile of the Hispanic foreign-born share distribution.

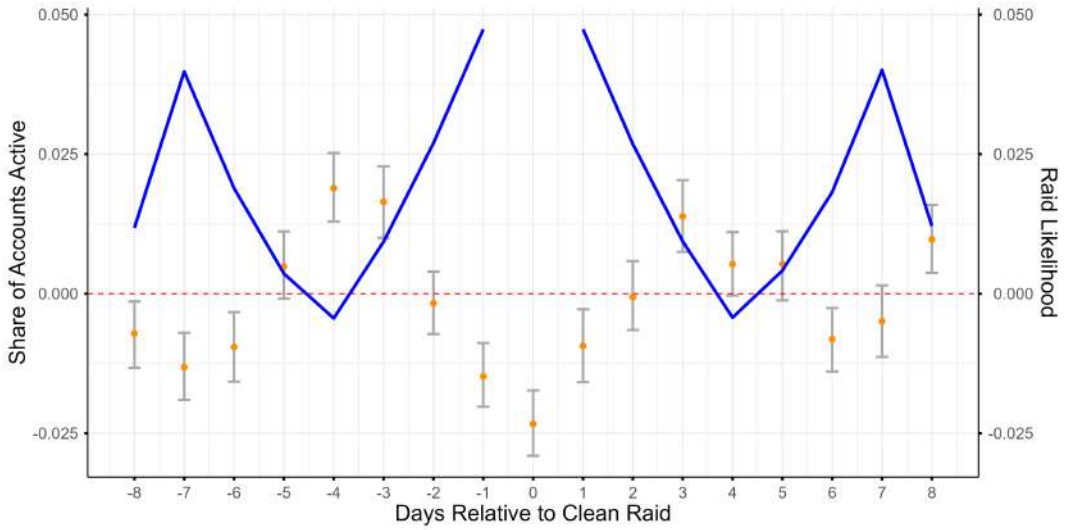
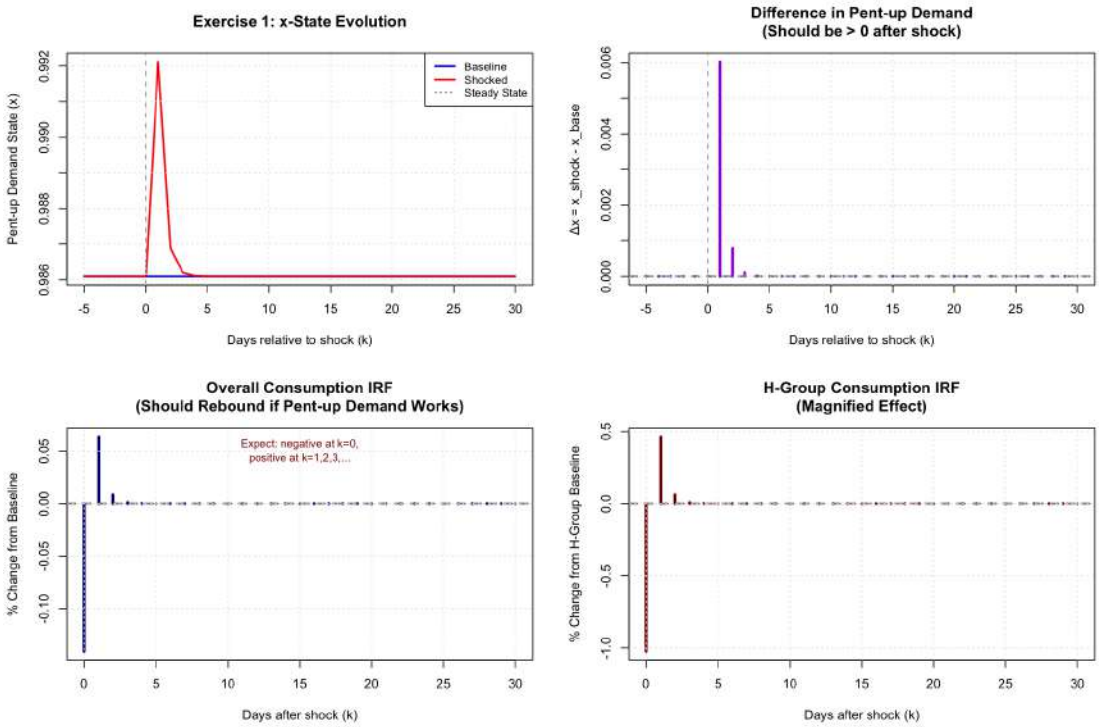
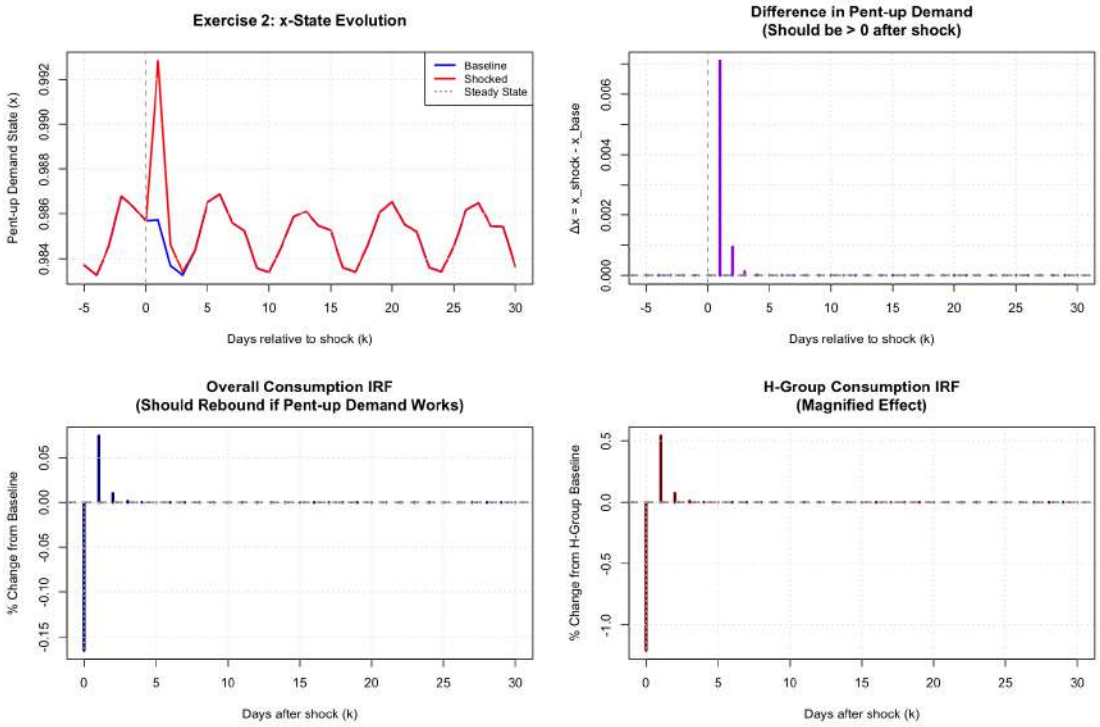


Figure 6: Clean Raids Event Study: Hispanic Foreign-born Share Interactions

Notes: Consumption coefficients (orange, left axis) overlaid with raid likelihood (blue, right axis) by days relative to clean raid. The event-study specification includes ZCTA and date fixed effects. The event study coefficients are those on the interaction between even time and the Hispanic foreign-born share. Error bars show 95 percent confidence intervals from 1,000 Bayesian bootstrap replications (Rubin, 1981) clustered at the ZCTA level.



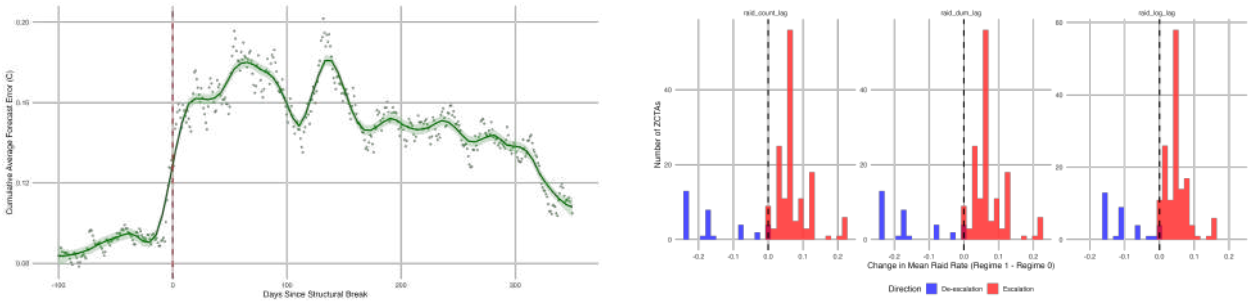
(a) Beliefs revert



(b) Beliefs evolve

Figure 7: Impulse Response to Belief Shocks

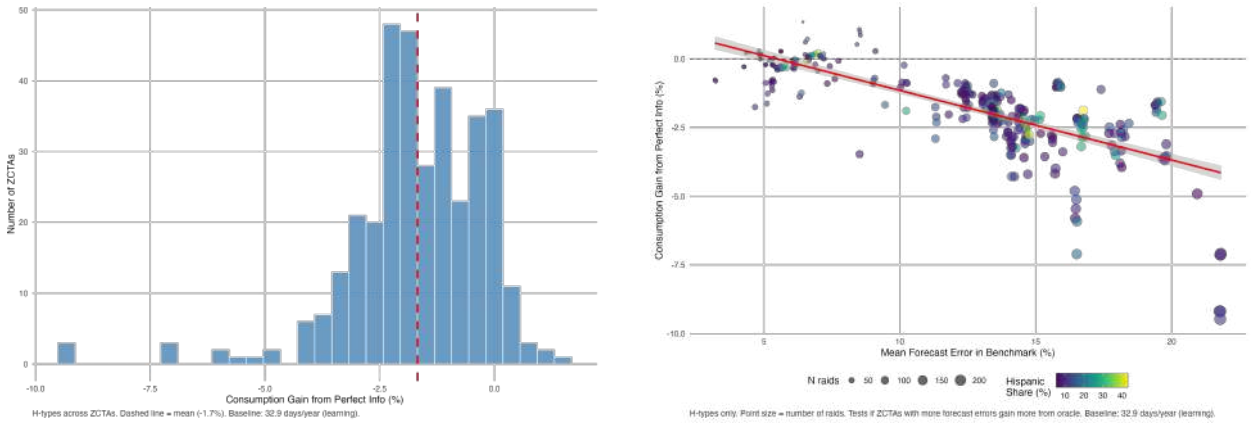
Notes: Impulse responses from a 15 percentage point belief shock from steady state. Panel (a) shows the response to a one-time shock to beliefs, where beliefs revert to baseline the next day. Panel (b) shows instead an exercise where, after the shock, beliefs are allowed to be updated over time. Each panel presents four subfigures. The top left presents the evolution of the pent-up demand stock. The top right presents the day-to-day changes in the pent-up demand stock. The bottom left panel presents the daily percent changes in economic activity overall. The bottom right panel presents the daily percent changes in economic activity among the Hispanic Foreign born group.



(a) Forecast errors around structural breaks **(b)** Distribution of regime changes

Figure 8: Structural Breaks as Model Validation

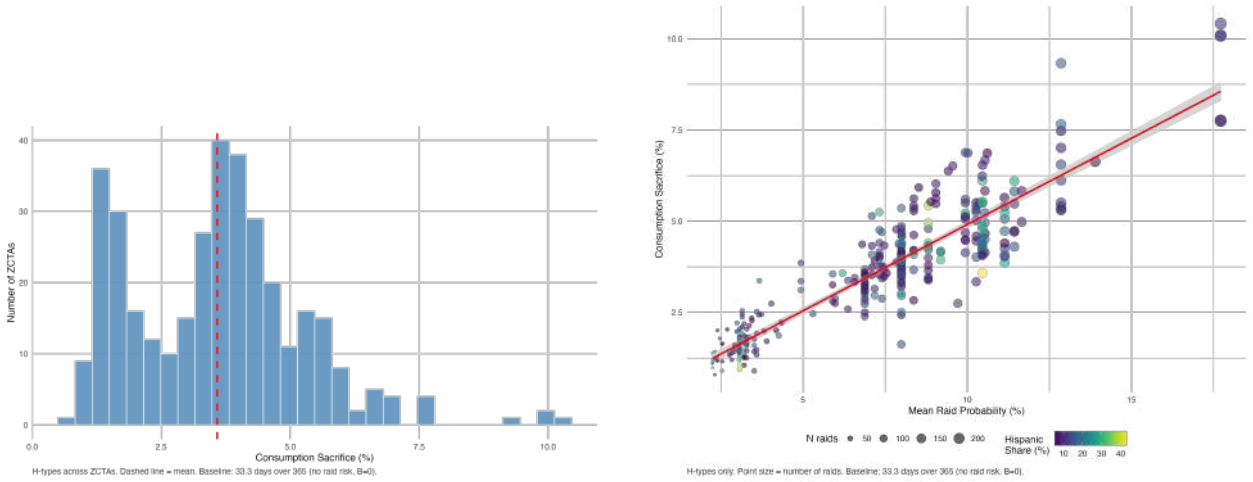
Notes: Panel (a) displays an event study of mean absolute forecast errors around detected structural breaks, with time normalized to zero at the break date. The shaded region shows the 95 percent confidence interval. Forecast errors spike sharply at the break and gradually decay as beliefs adjust to the new regime. Panel (b) shows the distribution of regime changes by direction: 82 percent of breaks represent escalations in enforcement intensity (red), with a mean increase of 6.5 percentage points in the daily raid probability.



(a) Distribution of consumption changes **(b)** Consumption change vs. forecast error

Figure 9: Perfect Information Counterfactual

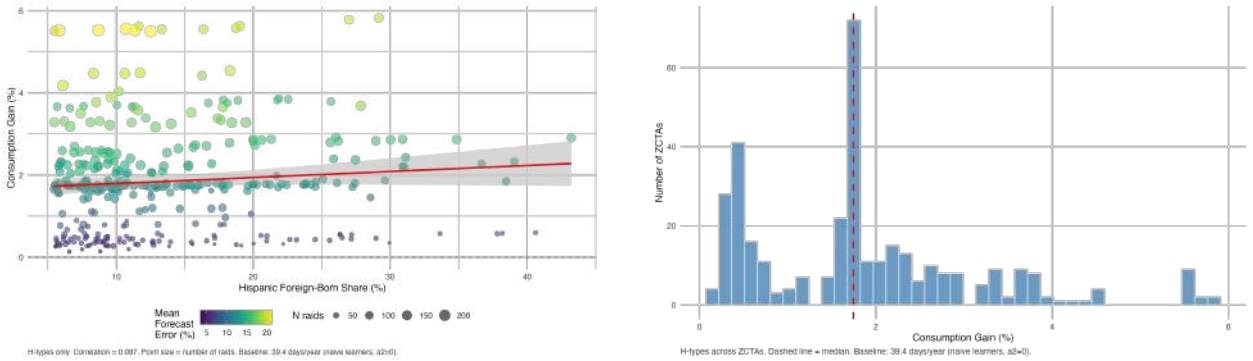
Notes: Panel (a) shows the distribution of consumption changes as a fraction of baseline consumption among the Hispanic foreign born under perfect information ($B_{zt} = r_{zt}$, $\bar{f}_{zt} = 0$) relative to the benchmark with Bayesian learning. The mean effect is -1.7 percent; 89 percent of ZCTAs experience consumption decreases. Panel (b) plots the corresponding consumption changes across ZCTAs against the model implied average cumulative forecast error. Dot colors represent the Hispanic foreign born share.



(a) Distribution of consumption sacrifice **(b)** Consumption sacrifice vs. raid probability

Figure 10: Willingness to Pay to Avoid Immigration Enforcement

Notes: Panel (a) shows the distribution of consumption sacrifice as a fraction of baseline consumption among the Hispanic foreign born comparing a scenario with no enforcement ($B = 0 \quad \forall t$) to one with known raid risk ($B = \bar{r}_{z,d}$). The mean sacrifice is 3.6 percent of consumption. Panel (b) plots the corresponding consumption changes across ZCTAs against the empirical average raid probabilities. Dot colors denote the Hispanic foreign born share.



(a) Gains vs. Hispanic foreign-born share **(b)** Distribution of consumption gains

Figure 11: The Value of Sophistication in Learning

Notes: Consumption gains from sophisticated learning ($\hat{a}_2 = 0.93$) relative to naive Bayesian updating ($a_2 = 0$). Panel (a) plots percent gains among the Hispanic foreign born against the Hispanic foreign-born share. Dot colors denote average forecast error. Panel (b) shows the distribution of percent gains among the Hispanic foreign born across ZCTAs. The mean gain is 1.85 percent (0.72 days/year).

Appendices

A Data Construction

A.1 Consumption

Data on consumer spending are drawn from the Facteus Debit and Credit Card Panel, a large-scale transaction dataset designed to be representative across income levels, geographic areas, and age groups in the United States. Facteus aggregates anonymized transaction data from a broad set of clients, including major debit, credit, and prepaid card processors, financial institutions, investment firms, and large retail corporations.

The dataset is among the largest proprietary card-transaction panels currently available, covering approximately 40 million unique cardholders, 20 million cards, and $1.3 trillion in consumer spending over the sample period. In aggregate, the panel captures roughly 4 percent of total U.S. consumer spending, which makes it particularly well suited for analysis at fine geographic resolutions, such as ZIP codes. The data were made available to us through a partnership between the Federal Reserve Bank of Chicago and Facteus.

The raw data are observed at the transaction level and include an anonymized account identifier, transaction date, ZIP code of residence of the account holder, transaction type (spend, deposit, ATM withdrawal, or load), and transaction value. We retain only accounts that are active both at the beginning and at the end of the observation window, defined as having at least one recorded transaction in both 2014 and 2018. This restriction mitigates bias arising from account entry and exit into the panel.

For the analysis, we aggregate transactions up to the ZCTA–day level. To ensure adequate and stable coverage across locations, we impose a "selection cone" that requires a minimum of 30 active accounts per ZIP code and a penetration rate between 3 and 25 accounts per 1,000 residents. This procedure excludes areas with insufficient data coverage as well as ZIP codes with implausibly high representation in the panel.

A.2 Geographic Mapping

The TRAC arrest records identify enforcement events at the county level, while our financial transaction data and demographic variables are measured at the ZIP Code Tabulation Area (ZCTA) level. Since ZCTAs do not nest cleanly within counties—many ZCTAs straddle county boundaries—we develop a two-step mapping procedure to assign county-level arrests to ZCTAs.

The first step maps county names in the TRAC data to county FIPS codes. The TRAC

records identify counties by name and state, but do not include standardized FIPS codes. We use a county crosswalk file developed for Ciancio and García-Jimeno (2024), which maps county names to FIPS codes and was constructed to handle variations in county naming conventions across data sources.

The second step maps counties to ZCTAs. We use the Census Bureau's ZCTA-to-county relationship file, which reports the share of each ZCTA's population residing in each overlapping county. We associate a ZCTA with a county if either (i) more than 2.5 percent of the ZCTA's population resides in that county, or (ii) the county is the only one the ZCTA overlaps with. Under this rule, ZCTAs may be associated with multiple counties—in our data, up to five.

For each ZCTA-date observation, we sum arrests across all associated counties. This approach treats a ZCTA as "exposed" to enforcement in any of its associated counties. For example, if ZCTA z is associated with counties c_1 and c_2, then arrests in ZCTA z on date t equals the sum of arrests in c_1 and c_2 on date t. This assignment rule implies that the same county-level arrest may contribute to the arrest count for multiple ZCTAs, which is appropriate given that residents of border-straddling ZCTAs may work, shop, or travel in any of their associated counties.

The mapping procedure yields 17,765 ZCTAs with at least one associated county appearing in the TRAC arrest records. After merging with Facteus transaction data and ACS demographics and applying sample restrictions (see Section A.3), the final analysis sample contains 12,865 ZCTAs.

A.3 Sample Selection

The analysis sample is constructed by merging three data sources: financial transaction data from Facteus, ICE interior arrest records from TRAC, and demographic characteristics from the 2018 American Community Survey 5-Year estimates. Table A1 documents the sample restrictions applied and their effects on sample size.

Step	Restriction	ZCTAs	ZCTA-Days
0	Facteus data (all zip codes)	39,648	53,088,672
1	Merge with Raids data	39,648	53,088,672
2	Merge with ACS data	30,177	40,407,003
3	Selection cone: accts $>$ 30, accts.pc $\in$ [3, 25)	12,865	17,226,235

Table A1: Sample Selection

Notes: Sample construction from Facteus financial transaction data, TRAC ICE arrest records, and 2018 ACS 5-Year estimates. ZCTA-Days computed as ZCTAs $\times$ 1,339 days (October 2014 – May 2018).

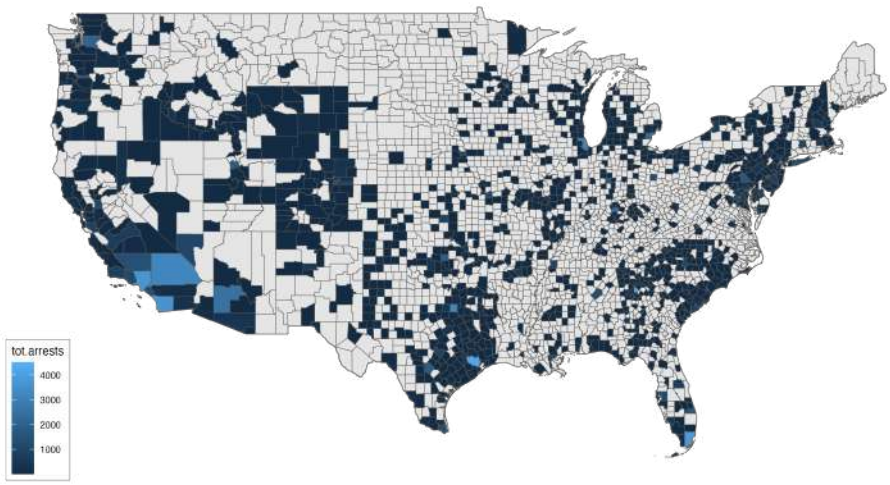
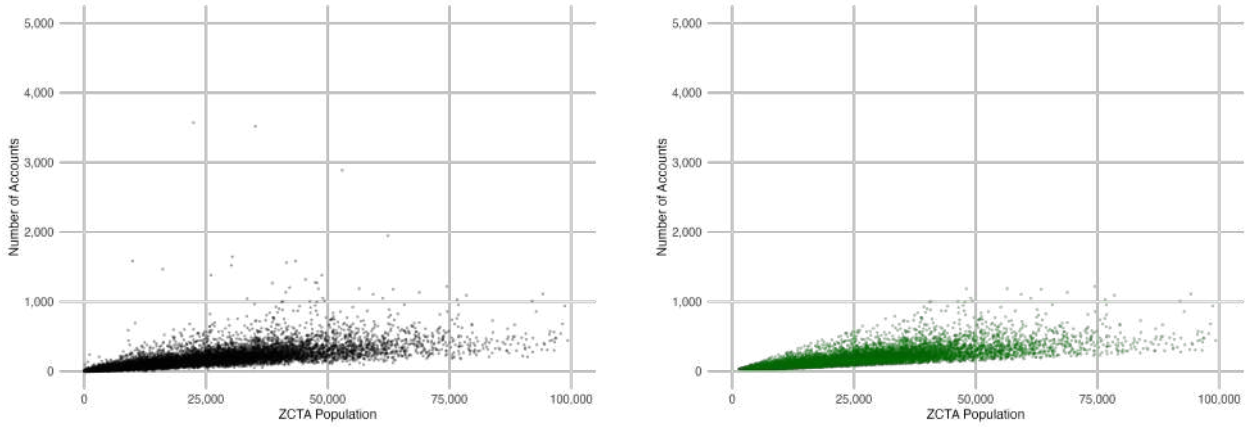


Figure A1: Geographic Distribution of ICE Interior Arrests, October 2014 – May 2018.
Notes: Total ICE interior raid arrests by county over the sample period. Brighter shading indicates higher arrest counts.

The starting point is the universe of ZIP codes appearing in the Facteus data during the sample period (October 2014 through May 2018). We first merge ICE arrest records, assigning zero arrests to ZIP codes with no recorded enforcement activity. We then perform an inner join with ACS demographic data, restricting to ZIP Code Tabulation Areas (ZCTAs) that appear in both datasets. This reduces the sample from 39,648 to 30,177 ZCTAs, reflecting the imperfect correspondence between commercial ZIP codes and Census-defined ZCTAs.

We impose one additional sample restriction: a "selection cone" requiring ZCTAs to have more than 30 accounts, at least 3 accounts per 1,000 population, and fewer than 25 accounts per 1,000 population. This ensures sufficient account coverage for meaningful activity measures while excluding ZCTAs with implausibly high or low penetration rates. Figure A2 illustrates the selection cone graphically.

The final sample contains 12,865 ZCTAs observed over 1,339 days, for a total of 17,226,235 ZCTA-day observations.



(a) All ZCTAs **(b)** Selection cone applied

Figure A2: Selection Cone for Sample Restriction

Notes: Scatter plots of ZCTA population against number of Facteus accounts. Left panel shows all ZCTAs after merging with ACS data. Right panel shows ZCTAs satisfying the selection cone restriction: more than 30 accounts, at least 3 accounts per 1,000 population, and fewer than 25 accounts per 1,000 population.

B Operation Crosscheck

Figure A3 displays the geographic distribution of ZCTAs included in the Operation Crosscheck event study sample. It includes 802 ZCTAs that experienced at least one arrest during the Crosscheck week (March 1–5, 2015) and had no arrests in the 10 days before or after.

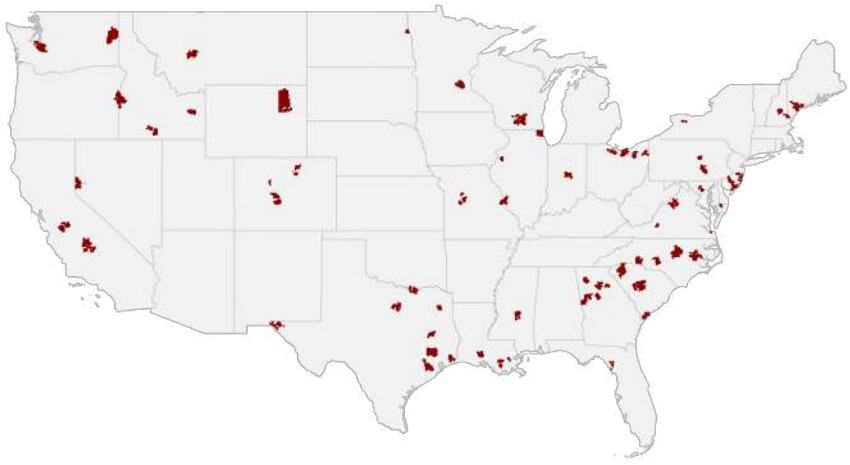


Figure A3: Geographic Distribution of Operation Crosscheck Event Study Sample

Notes: Map shows the 802 ZCTAs included in the Crosscheck event study sample. ZCTAs experienced at least one arrest during Operation Crosscheck (March 1–5, 2015) with no arrests in the 10 days before or after.

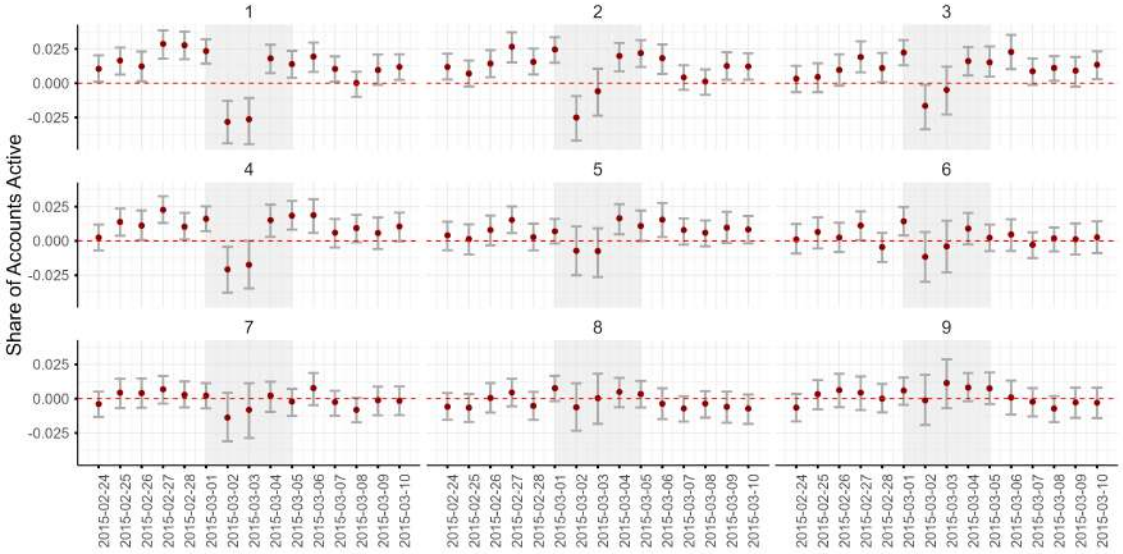


Figure A4: Operation Crosscheck Event Study: Effects by Hispanic Foreign-Born Decile

Notes: Event study coefficients by Hispanic foreign-born share decile. Decile 1 corresponds to ZCTAs with the highest Hispanic foreign-born shares; Decile 10 corresponds to the lowest. Sample and estimation as in Figure 2. Error bars show 95 percent confidence intervals from 1,000 Bayesian bootstrap replications (Rubin, 1981) clustered at the ZCTA level.

C Event Studies Robustness

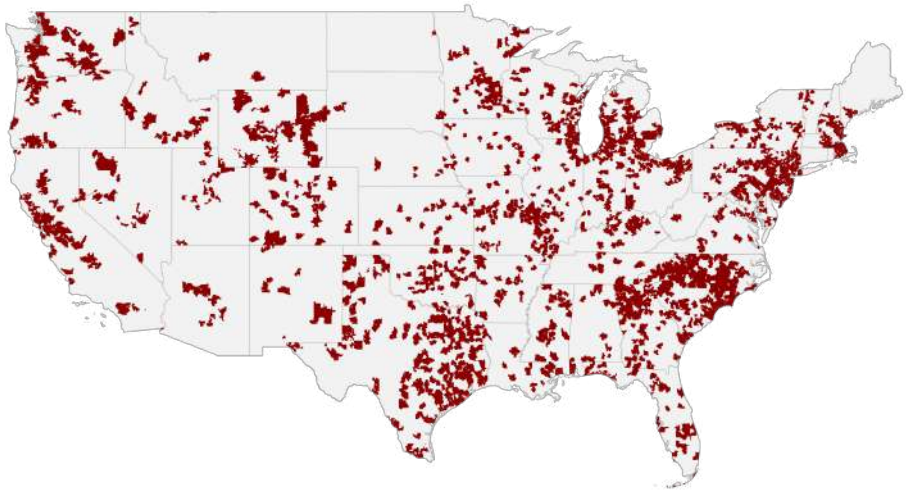


Figure A5: Geographic Distribution of Event Study Sample
Notes: Map shows the 5,967 ZCTAs included in the event study second stage. These ZCTAs experienced at least one "clean" raid event (no other raids within ±8 days) during the sample period.

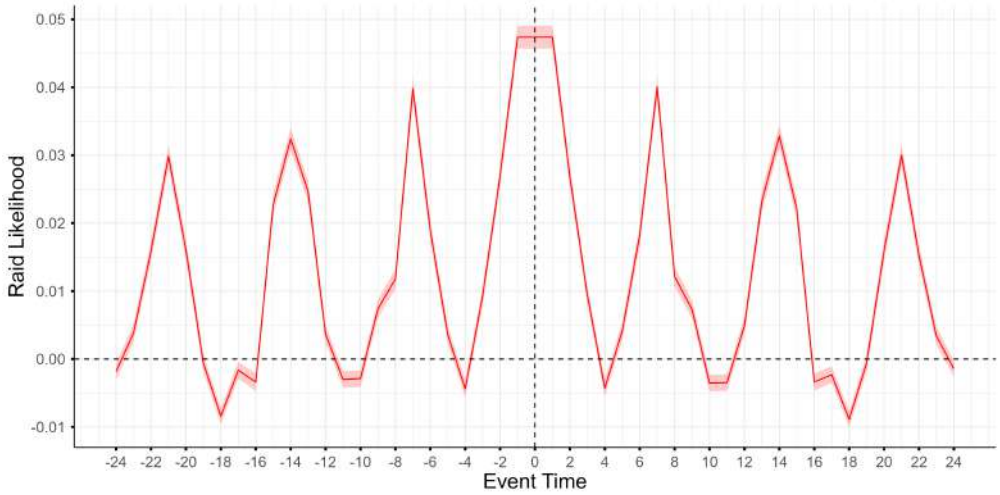


Figure A6: Raid Likelihood from Autocorrelation Regression
Notes: Coefficients from regressing the raid indicator on 24 leads and lags of itself, with ZCTA and date fixed effects. The seven-day periodicity reflects ICE's tendency to conduct raids on preferred weekdays. Sample restricted to ZCTAs in the event study estimation sample.

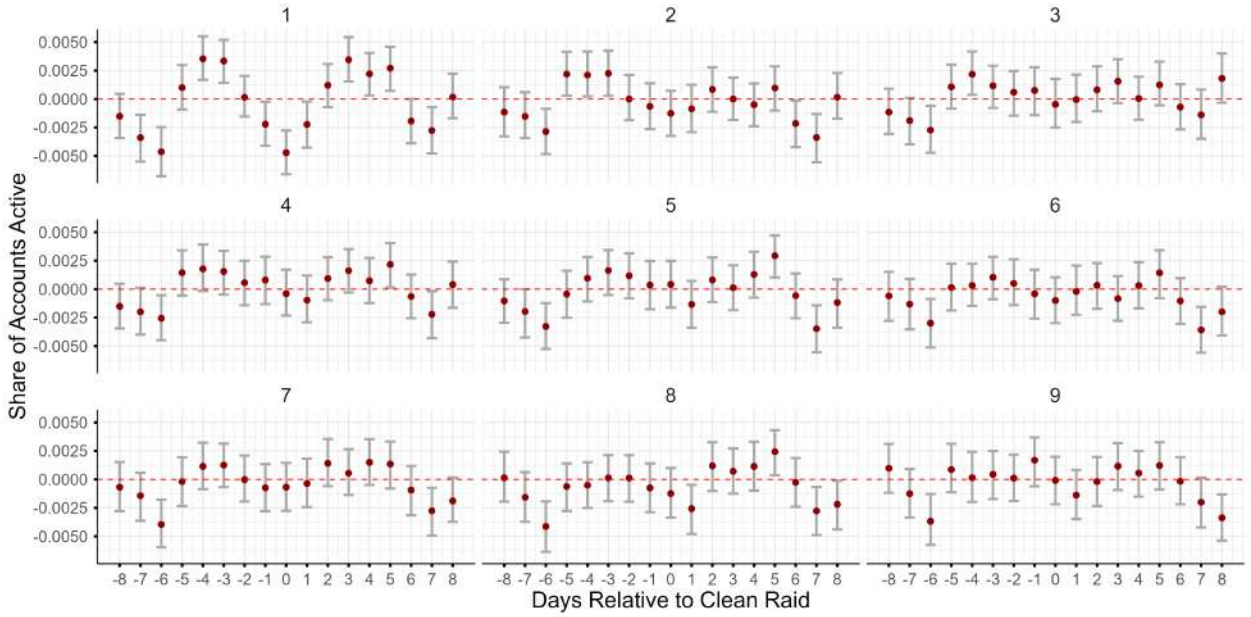


Figure A7: Clean Raids Event Study: Effects by Hispanic Foreign-Born Share Decile

Notes: Event study coefficients by decile of Hispanic foreign-born share. Decile 1 corresponds to ZCTAs with the highest Hispanic foreign-born shares; Decile 9 corresponds to the lowest shown. The seven-day cyclical pattern is strongest and most statistically significant in the top deciles. Effects attenuate progressively toward lower deciles. Sample includes 33,091 clean raid events across 5,967 ZCTAs. Error bars show 95 percent confidence intervals from 1,000 Bayesian bootstrap replications (Rubin, 1981) clustered at the ZCTA level.

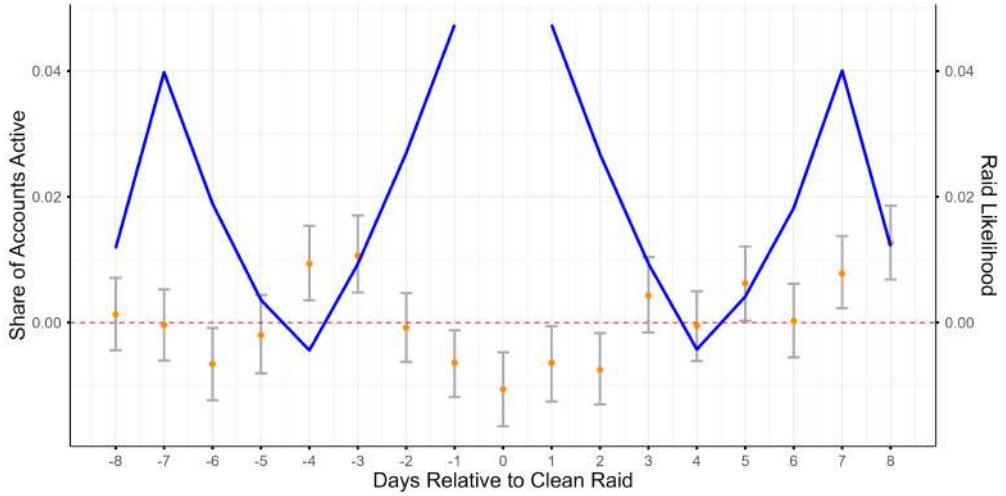


Figure A8: Clean Raids Event Study: Hispanic Foreign-born Share Interactions including ZCTA × Weekday Fixed Effects

Notes: Consumption coefficients (orange, left axis) overlaid with raid likelihood (blue, right axis) by days relative to clean raid. The event-study specification includes ZCTA × weekday and date fixed effects. The event study coefficients are those on the interaction between even time and the Hispanic foreign-born share. Error bars show 95 percent confidence intervals from 1,000 Bayesian bootstrap replications (Rubin, 1981) clustered at the ZCTA level.

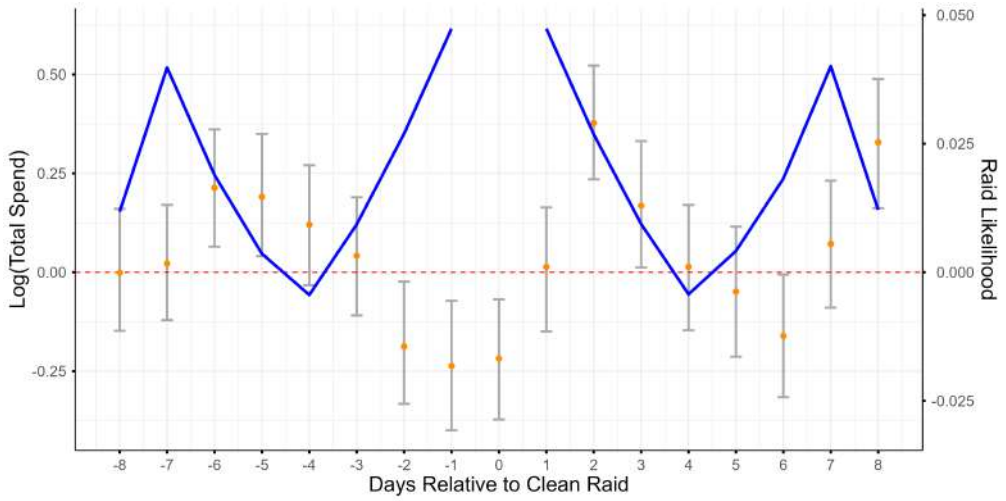


Figure A9: Event Study Overlay: Log Total Spending

Notes: Consumption coefficients (orange, left axis) overlaid with raid likelihood (blue, right axis) by days relative to clean raid. Outcome variable is log total spending. The event study coefficients are those on the interaction between even time and the Hispanic foreign-born share. Error bars show 95 percent confidence intervals from 1,000 Bayesian bootstrap replications (Rubin, 1981) clustered at the ZCTA level.

D Measurement Error Derivations

This appendix provides the full derivation of the OLS bias result and sign condition stated in Section 3.5.

D.1 Setup and Notation

Let $r_{zt} \in \{0, 1\}$ denote whether a raid occurs in ZCTA z on date t, and let $r^e_{zt} \equiv \mathbb{E}[r_{zt}|\mathcal{I}_{zt}]$ denote agents' subjective belief about raid likelihood given their information set. Define the prediction error as $\tilde{\eta}_{zt} \equiv r_{zt} - r^e_{zt}$. Let $\varphi^* \equiv \mathbb{P}(r_{zt} = 1)$ denote the true (unconditional) raid probability.

The model is

$$y_{zt} = \alpha_z + \delta_t + \beta r^e_{zt} + \epsilon_{zt}$$

Using r_{zt} as a proxy for r^e_{zt} yields

$$y_{zt} = \alpha_z + \delta_t + \beta r_{zt} + (\epsilon_{zt} - \beta\tilde{\eta}_{zt})$$

Assuming $\text{cov}(r_{zt}, \epsilon_{zt}) = 0$ conditional on fixed effects, the OLS estimand is

$$\beta_{OLS} = \beta + \frac{\text{cov}(r_{zt}, \epsilon_{zt} - \beta\tilde{\eta}_{zt})}{\text{var}(r_{zt})} = \beta \left[1 - \frac{\text{cov}(r_{zt}, \tilde{\eta}_{zt})}{\text{var}(r_{zt})}\right]$$

We now compute $\text{cov}(r_{zt}, \tilde{\eta}_{zt})$. First:

$$\begin{aligned}\mathbb{E}[r_{zt}\tilde{\eta}_{zt}] &= \mathbb{E}\left[\mathbb{E}[r_{zt}\tilde{\eta}_{zt}|r_{zt}]\right] \\ &= \mathbb{P}(r_{zt}=1)\mathbb{E}[\tilde{\eta}_{zt}|r_{zt}=1] \\ &= \varphi^*\mathbb{E}[r_{zt}-r^e_{zt}|r_{zt}=1] \\ &= \varphi^*(1-\mathbb{E}[r^e_{zt}|r_{zt}=1])\end{aligned}$$

Next:

$$\begin{aligned}\mathbb{E}[\tilde{\eta}_{zt}] &= \mathbb{E}\left[\mathbb{E}[\tilde{\eta}_{zt}|r_{zt}]\right] \\ &= \varphi^*\mathbb{E}[r_{zt}-r^e_{zt}|r_{zt}=1] + (1-\varphi^*)\mathbb{E}[r_{zt}-r^e_{zt}|r_{zt}=0] \\ &= \varphi^*(1-\mathbb{E}[r^e_{zt}|r_{zt}=1]) - (1-\varphi^*)(\mathbb{E}[r^e_{zt}|r_{zt}=0])\end{aligned}$$

Therefore:

$$\begin{aligned}\text{cov}(r_{zt}, \tilde{\eta}_{zt}) &= \mathbb{E}[r_{zt}\tilde{\eta}_{zt}] - \mathbb{E}[r_{zt}]\mathbb{E}[\tilde{\eta}_{zt}] \\ &= \varphi^*(1-\mathbb{E}[r^e_{zt}|r_{zt}=1]) \\ &\quad - \varphi^*\left\{\varphi^*(1-\mathbb{E}[r^e_{zt}|r_{zt}=1]) - (1-\varphi^*)(\mathbb{E}[r^e_{zt}|r_{zt}=0])\right\} \\ &= \varphi^*(1-\varphi^*)\left\{1-\mathbb{E}[r^e_{zt}|r_{zt}=1]+\mathbb{E}[r^e_{zt}|r_{zt}=0]\right\}\end{aligned}$$

Since $\text{var}(r_{zt}) = \varphi^*(1-\varphi^*)$ for a Bernoulli random variable:

$$\begin{aligned}\beta_{OLS} &= \beta\left[1 - \frac{\varphi^*(1-\varphi^*)\left\{1-\mathbb{E}[r^e_{zt}|r_{zt}=1]+\mathbb{E}[r^e_{zt}|r_{zt}=0]\right\}}{\varphi^*(1-\varphi^*)}\right] \\ &= \beta\left[1-\left\{1-\mathbb{E}[r^e_{zt}|r_{zt}=1]+\mathbb{E}[r^e_{zt}|r_{zt}=0]\right\}\right] \\ &= \beta\left[\mathbb{E}[r^e_{zt}|r_{zt}=1]-\mathbb{E}[r^e_{zt}|r_{zt}=0]\right]\end{aligned}$$

This is the OLS bias result stated in equation (8).

D.2 Condition for Correct Sign

When does the OLS estimand have the same sign as the true β? Using Bayes' rule, we can express the conditional expectations as:

$$\mathbb{E}[r^e_{zt}|r_{zt}=r] = \int r^e_{zt}\frac{\mathbb{P}(r_{zt}=r|r^e_{zt})f(r^e_{zt})}{\mathbb{P}(r_{zt}=r)}dr^e_{zt}$$

Let $P(r^e_{zt}) \equiv \mathbb{P}(r_{zt}=1|r^e_{zt})$ denote the true raid probability given beliefs. After some

algebra (expanding and simplifying), the difference $\mathbb{E}[r^e_{zt}|r_{zt} = 1] - \mathbb{E}[r^e_{zt}|r_{zt} = 0]$ has the same sign as β if and only if:

$$\text{cov}(r^e_{zt}, P(r^e_{zt})) \geq 0$$

This condition says that the OLS has the correct sign as long as beliefs and true raid probabilities are positively correlated: when agents believe raids are more likely, actual raid probabilities should also be higher. This condition is satisfied whenever learning is taking place, since learning implies that high beliefs correspond to genuinely high-risk periods. Notice that when beliefs contain no information about raid likelihood (i.e., $P(r^e_{zt}) = \varphi^*$ for all r^e_{zt}), then $\text{cov}(r^e_{zt}, P(r^e_{zt})) = 0$ and the OLS estimand is zero. Thus, OLS can be interpreted as a test of whether beliefs contain information about the raid process.

E Regime Changes

Figure A10 presents the distribution of structural break dates across all ZCTAs with breaks.

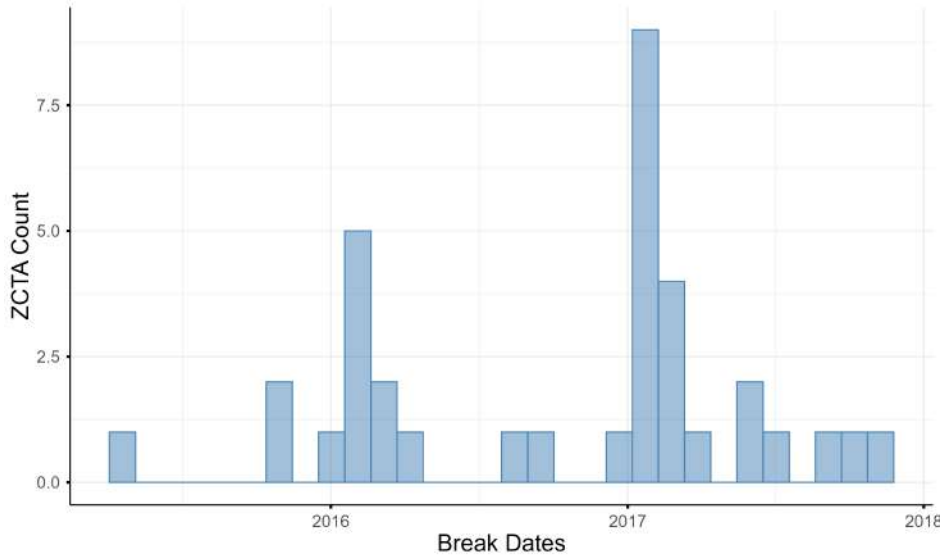


Figure A10: Distribution of Detected Structural Breaks in ICE Enforcement

*Notes:*Histogram of structural break dates detected using Bai and Perron (2003) methodology. 513 breaks across 439 ZCTAs. Sample restricted to top-quintile Hispanic foreign-born ZCTAs with at least 191 raid days. Breaks graded by agreement across three specifications (raid dummy, raid count, log raid count with weekday dummies and lag).

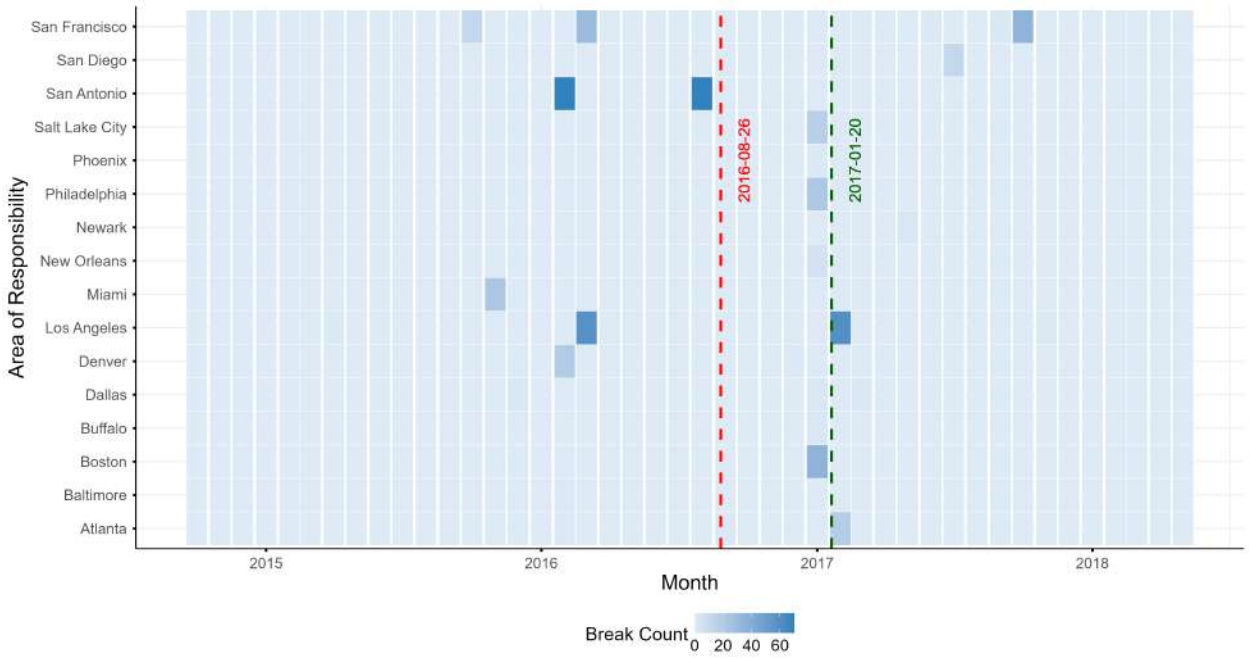
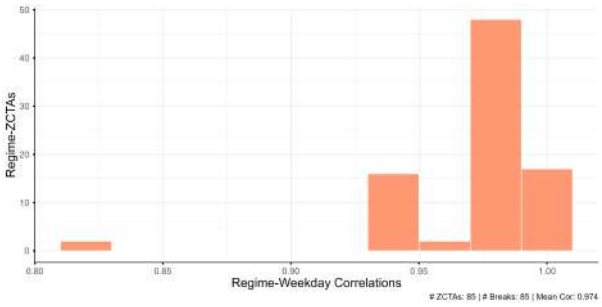
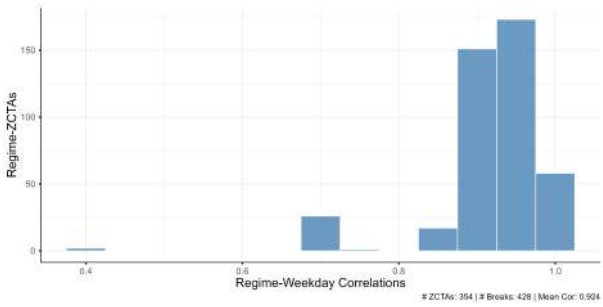


Figure A11: Timing of Structural Breaks by ICE District

Notes: Distribution of structural break dates by ICE Area of Responsibility (district). Structural break dates detected using Bai and Perron (2003) methodology. Sample restricted to top-quintile Hispanic foreign-born ZCTAs with at least 191 raid days. Breaks graded by agreement across three specifications (raid dummy, raid count, log raid count with weekday dummies and lag). Districts ordered by median break date. Red dashed line indicates the overall median break date. Green dashed line indicates President Trump's first inauguration. N = 513 breaks across 439 ZCTAs.



(a) Breaks within a month of Trump Inauguration



(b) Breaks outside a month of Trump Inauguration

Figure A12: Correlation of Weekday Betas Across Regimes

Notes: Distribution of correlations between weekday coefficients in regime 0 (before break) and regime 1 (after break). Structural break dates detected using Bai and Perron (2003) methodology. Sample restricted to top-quintile Hispanic foreign-born ZCTAs with at least 191 raid days. Breaks graded by agreement across three specifications (raid dummy, raid count, log raid count with weekday dummies and lag). Panel (a) includes ZCTAs with detected breaks within one month of Trump inauguration (N = 85 ZCTAs); panel (b) includes breaks more than one month from inauguration (N = 354 ZCTAs). High correlation indicates pure intensity shift (same days remain high/low risk); low correlation indicates re-shuffling of weekday patterns.

F Estimation Details

This appendix provides technical details on the estimation procedure summarized in Section 5.

F.1 Value Function Solution

For a given parameter vector $\boldsymbol{\theta} = (\psi, \delta, \alpha_{11}, \alpha_{12})$, we solve for the stationary value function $V(x)$ by value iteration over a discrete grid.

Grid construction. We discretize the pent-up state space $x \in [x_{\min}, x_{\max}]$ using a grid of 401 equally-spaced points. The upper bound $x_{\max} = 25$ ensures the stationary distribution of x has negligible mass near the boundary.

Iteration. Starting from an initial guess $V^{\tau(0)}(x) = 0$, we iterate on the Bellman equation

$$V^{\tau(k+1)}(x) = \log\Big(\exp\big(c(x) + \beta V^{\tau(k)}((1-\delta)x + \kappa)\big) + \exp\big(c(x) + \beta V^{\tau(k)}((1-\delta)x)\big)\Big), \quad (20)$$

where $c(x) = -(\psi/2)x^2$ is the flow cost. We compute between-grid values by linear interpolation. The iteration continues until $\|V^{\tau(k+1)} - V^{\tau(k)}\| < 10^{-8}$.

Dynamic incentive. Once convergence is achieved, we compute the dynamic incentive term

$$\Delta(x)^{\tau} \equiv \beta\big[V^{\tau}((1-\delta)x) - V^{\tau}((1-\delta)x + \kappa)\big]. \quad (21)$$

This quantity enters the probabilities and captures the option value of going out today versus accumulating additional pent-up demand.

Figure A13 displays the estimated value function and dynamic incentive term at the estimated parameter values.

F.2 Belief Construction

Prior initialization. We estimate the prior from a probit regression on the first 105 days (burn-in period):

$$r_{zt} = \Phi\big(\gamma_{d(t),r_{t-1}} + \beta^{h}_{d(t),r_{t-1}} h_z + \beta^{y}_{d(t),r_{t-1}} y_z + \eta_{\mathrm{AOR}(z)}\big), \quad (22)$$

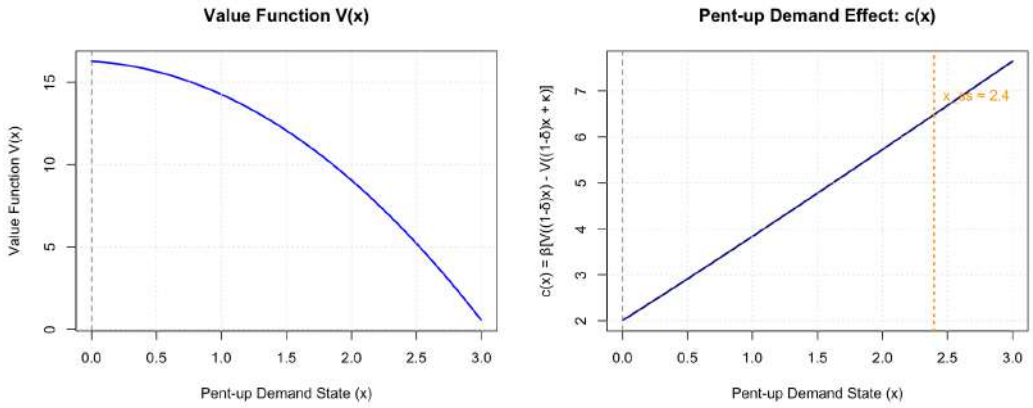


Figure A13: Estimated Value Function and Dynamic Incentive

Notes: Left panel shows the value function $V(x)$ as a function of the pent-up demand state x. Right panel shows the dynamic incentive term $\Delta(x) = \beta[V((1-\delta)x) - V((1-\delta)x + \kappa)]$, which captures the option value of going out today versus accumulating additional pent-up demand. Both computed at the estimated parameter values $\hat{\psi} = 2.23$ and $\hat{\delta} = 0.42$.

where $\gamma_{d(t),r_{t-1}}$ are weekday $\times$ previous-day-raid intercepts, $\beta^h_{d(t),r_{t-1}}$ and $\beta^y_{d(t),r_{t-1}}$ are weekday $\times$ previous-day-raid slopes on Hispanic foreign-born share h_z and log median household income y_z, respectively, and $\eta_{\text{AOR}(z)}$ are ICE Area of Responsibility fixed effects. The predicted probability $\hat{q}_{z,d,k}$ initializes the Beta prior mean, with $a^0 + b^0 = 10$. Table A2 reports the probit estimates.

Bayesian updating. For each ZCTA z, weekday d, and lag state $k \in \{0, 1\}$, we compute the Beta parameter updates according to equation (17).

Cumulative forecast error. We compute the running average absolute forecast error

$$\bar{f}_{zt} = \frac{1}{t - t_0} \sum_{s=t_0}^{t-1} \omega_s |r_{zs} - B_{zs}|, \tag{23}$$

where t_0 is the end of the burn-in period. This measure enters the heterogeneous belief coefficient $\alpha_1(\bar{f}) = \alpha_{11} + \alpha_{12}\bar{f}$.

F.3 Mixture Model Inversion

Given beliefs $\{r^e_{zt}\}$, the H-type and N-type choice probabilities are

$$A^H_{zt} = \Lambda\big(\alpha_1(\bar{f}_{zt}) r^e_{zt} + \Delta(x^H_{zt}) + s_{zt}\big),$$
$$A^N_{zt} = \Lambda\big(\Delta(x^N_{zt}) + s_{zt}\big),$$

	$r_{t-1} = 0$		$r_{t-1} = 1$	
Weekday	Coef.	SE	Coef.	SE
Panel A: Baseline Intercepts				
Monday	-1.01	(2.44)	-3.86	(1.60)
Tuesday	-2.41	(2.49)	0.13	(5.56)
Wednesday	0.01	(2.92)	-1.21	(2.99)
Thursday	0.59	(2.63)	-3.19	(3.89)
Friday	(ref.)		-9.10	(4.13)
Saturday	-14.82	(6.19)	-5.44	(1.62)
Sunday	3.10	(1.68)	-5.40	(2.27)
Panel B: Hispanic Foreign-Born Share Interactions				
Monday	-1.21	(0.79)	0.00	(0.00)
Tuesday	0.29	(0.71)	0.59	(2.26)
Wednesday	0.05	(0.77)	1.98	(1.43)
Thursday	-0.48	(0.74)	2.46	(1.86)
Friday	1.86	(0.69)	4.02	(0.80)
Saturday	3.79	(1.77)	-0.43	(0.62)
Sunday	-2.62	(0.52)	0.27	(0.38)
Panel C: Log Median Household Income Interactions				
Monday	0.08	(0.19)	0.00	(0.00)
Tuesday	0.23	(0.16)	0.04	(0.46)
Wednesday	-0.00	(0.24)	0.16	(0.27)
Thursday	-0.06	(0.20)	0.33	(0.34)
Friday	-0.03	(0.15)	0.80	(0.43)
Saturday	1.21	(0.54)	0.13	(0.10)
Sunday	-0.41	(0.10)	0.13	(0.17)
AOR Fixed Effects	Yes (16 districts)			
Observations	36,608			
Log-likelihood	-6158.4			

Table A2: Probit Regression for Prior Belief Initialization

Notes: Probit regression of raid indicator on weekday × previous-day-raid fixed effects interacted with ZCTA demographics. Sample restricted to first 105 days (burn-in period) for ZCTAs in the top quintile of Hispanic foreign-born share with at least 30 raids. $r_{t-1} = 0$ indicates no raid on the previous day; $r_{t-1} = 1$ indicates a raid occurred. Standard errors in parentheses, two-way clustered by ZCTA and date.

where $\Lambda(\cdot)$ is the logistic CDF and s_{zt} is the structural shock. The observed economic activity rate is the mixture

$$A_{zt} = h_z \cdot A_{zt}^H + (1 - h_z) \cdot A_{zt}^N.$$

Inversion procedure. Because s_{zt} enters both $\Lambda(\cdot)$ terms additively and Λ is strictly increasing, A_{zt} is strictly increasing in s_{zt} with range $(0, 1)$. Thus, any observed A_{zt} admits a unique s_{zt} that rationalizes it. For each observation (z, t), we seek the shock s_{zt} that rationalizes the observed A_{zt}, which we find solving the fixed-point problem with a simple bisection algorithm.

We proceed as follows:

1. Initialize pent-up states: $x_{z,t_0}^H = x_{z,t_0}^N = 0$.
2. For each $t = t_0 + 1, \ldots, T$:
   (a) Given current states (x_{zt}^H, x_{zt}^N), use bisection to find $s_{zt} \in [-30, 30]$ such that the mixture probability equals A_{zt}.
   (b) Update states:

$$x_{z,t+1}^\tau = (1 - \delta)x_{zt}^\tau + \kappa - p_{zt}^\tau \tag{24}$$

3. Output the sequence of recovered shocks $\{s_{zt}\}$.

The bounds $s \in [-30, 30]$ ensure numerical stability; in practice, recovered shocks rarely approach these bounds.

F.4 Fixed Effects Concentration

The structural shocks s_{zt} should be orthogonal to the fixed effects structure if the model is correctly specified. We exploit this by concentrating out the fixed effects through regression. This is computationally convenient because our model has 1339-105 = 1234 time fixed effects, and $352 \times 7 = 2464$ ZCTA-by-weekday fixed effects.

Regression specification. We regress the recovered shocks on date and ZCTA-by-weekday fixed effects:

$$s_{zt} = \xi_t + \zeta_{z,d(t)} + \varepsilon_{zt}. \tag{25}$$

The date fixed effects ξ_t absorb aggregate shocks (holidays, macroeconomic conditions). The ZCTA-by-weekday fixed effects $\zeta_{z,d(t)}$ absorb persistent within-week consumption patterns.

Objective function. The estimation objective is the residual sum of squares:

$$Q(\boldsymbol{\theta}) = \sum_{z,t} \hat{\varepsilon}_{zt}^2 = \sum_{z,t} \big(s_{zt}(\boldsymbol{\theta}) - \hat{\xi}_t - \hat{\zeta}_{z,d(t)}\big)^2. \tag{26}$$

F.5 Optimization

Algorithm. We minimize $Q(\boldsymbol{\theta})$ using L-BFGS-B (Byrd et al., 1995), a quasi-Newton method that accommodates box constraints. The parameter bounds are:

- $\psi \in [0.0001, 10]$ (pent-up cost)
- $\delta \in [0.02, 0.98]$ (depreciation rate)
- $\alpha_{11} \in [-20, 20]$ (belief effect)
- $\alpha_{12} \in [-20, 20]$ (interaction coefficient)

To guard against local optima, we employ a multi-start optimization with 24 random initial points drawn uniformly from the parameter bounds.

Figure A14 displays the profiled objective function for each parameter, holding the others fixed at their estimated values. The profiles are well-behaved with clear global minima, confirming that our estimates correspond to a unique optimum.

F.6 Standard Errors

We compute standard errors in two steps.

Numerical Hessian. We approximate the Hessian matrix $H = \nabla^2 Q(\hat{\boldsymbol{\theta}})$ using central finite differences with step size 10^{-4}.

HAC adjustment. Because the recovered shocks may exhibit serial correlation and heteroskedasticity, we compute heteroskedasticity and autocorrelation-consistent (HAC) standard errors using the Newey-West estimator (Newey and West, 1987). We report results across a range of lag specifications to assess robustness.

The asymptotic variance is

$$\widehat{\mathrm{Var}}(\hat{\boldsymbol{\theta}}) = H^{-1} \hat{\Omega} H^{-1}, \tag{27}$$

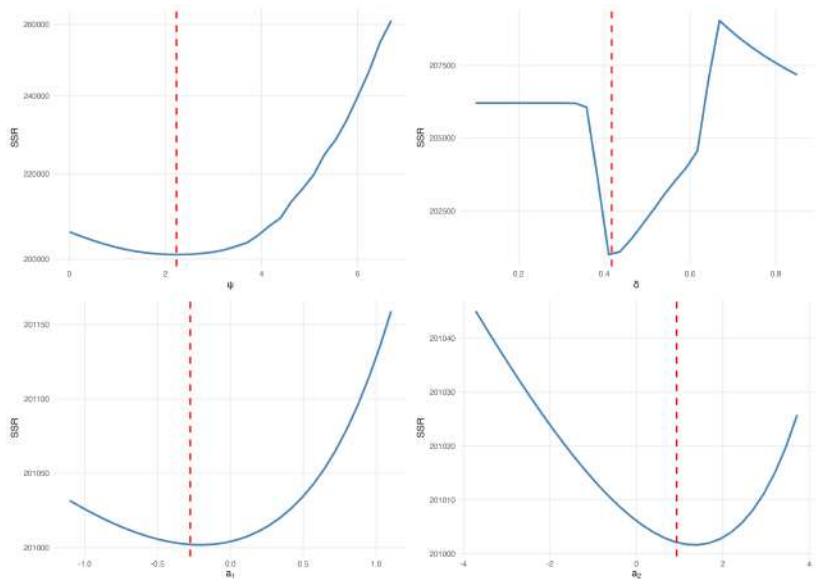


Figure A14: Profile Likelihood for Structural Parameters

Notes: Each panel shows the objective function $Q(\boldsymbol{\theta})$ as a function of one parameter, holding the remaining parameters fixed at their estimated values. Vertical dashed lines indicate the parameter estimates. The well-behaved profiles with clear global minima indicate that our estimates correspond to a unique optimum.

where $\hat{\Omega}$ is the HAC-adjusted outer product of the gradient.

Figure A15 displays the autocorrelation function of the recovered residuals, showing persistence at short lags that motivates the HAC adjustment.

F.7 Estimated Fixed Effects

Figure A16 displays the estimated fixed effects from the structural estimation.

F.8 Implied Beliefs: Sample Averages and ZCTA Examples

Figure A18 illustrates the implied beliefs from the structural model. Panel (a) shows sample-wide average beliefs by day of week over time. Panels (b) and (c) show belief evolution for two example ZCTAs, illustrating heterogeneity in belief dynamics across communities.

F.9 Impulse Responses

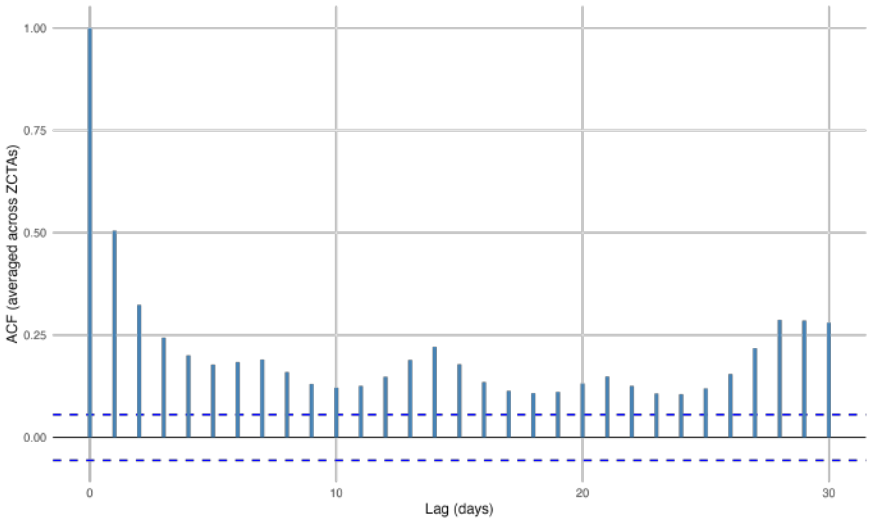
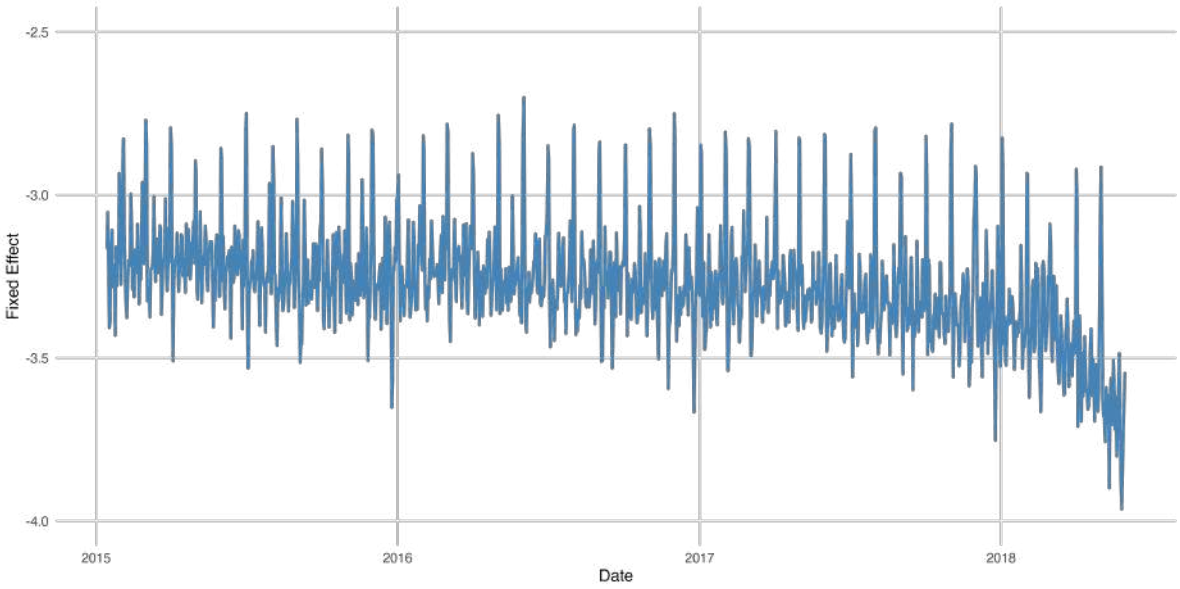
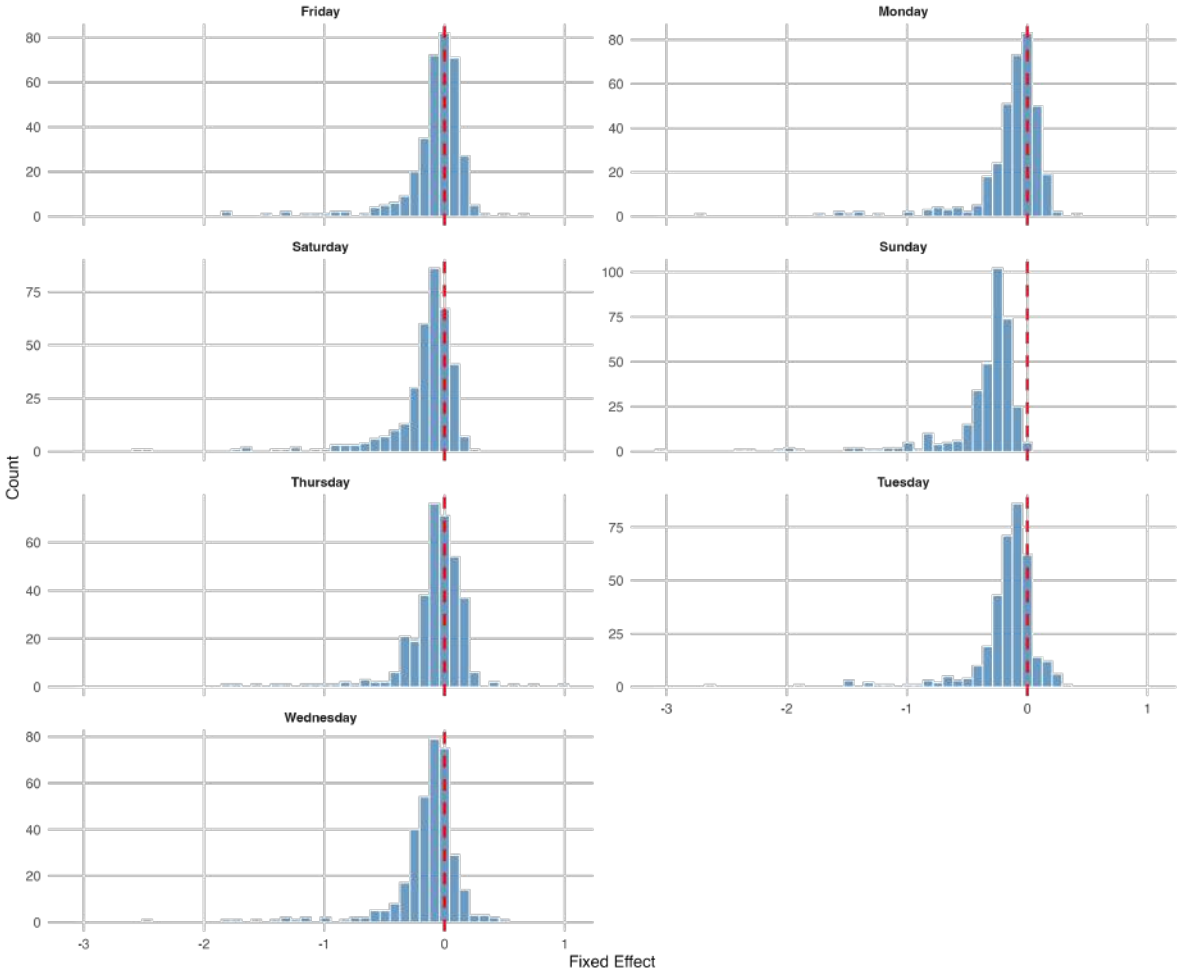


Figure A15: Autocorrelation Function of Residuals

Notes: Autocorrelation function of the recovered structural shocks $\hat{s}_{zt}$ from the mixture model inversion. The persistence at short lags motivates the use of Newey-West HAC standard errors. Dashed lines indicate 95 percent confidence bands under the null of no autocorrelation.



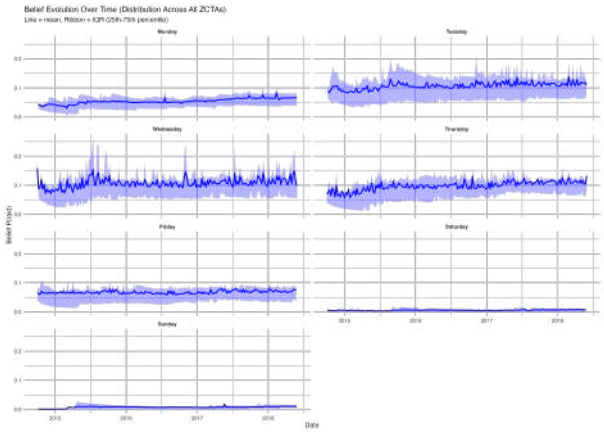
(a) Date fixed effects over time



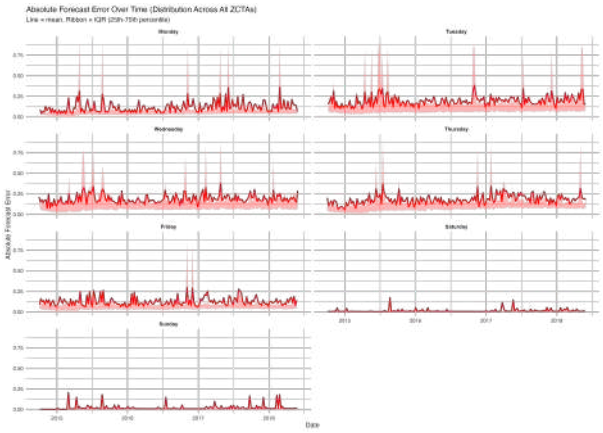
(b) Distribution of ZCTA-by-weekday fixed effects

Figure A16: Estimated Fixed Effects from Structural Model

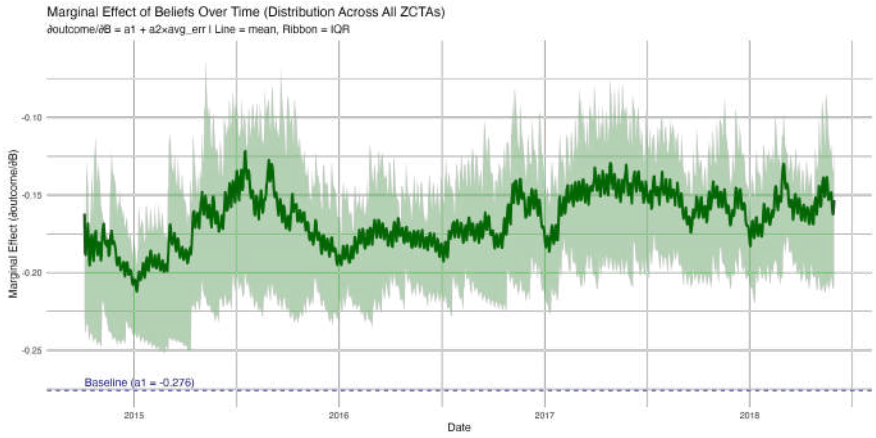
Notes: Panel (a) shows the estimated date fixed effects $\hat{\lambda}_t$ over the sample period, capturing aggregate shocks to economic activity (holidays, macroeconomic conditions). Panel (b) shows the distribution of ZCTA-by-weekday fixed effects $\hat{\mu}_{z,d}$, which absorb persistent within-week consumption patterns at the community level. The estimation includes 1,234 date fixed effects and 2,464 ZCTA-by-weekday fixed effects.



(a) Distribution of beliefs by weekday



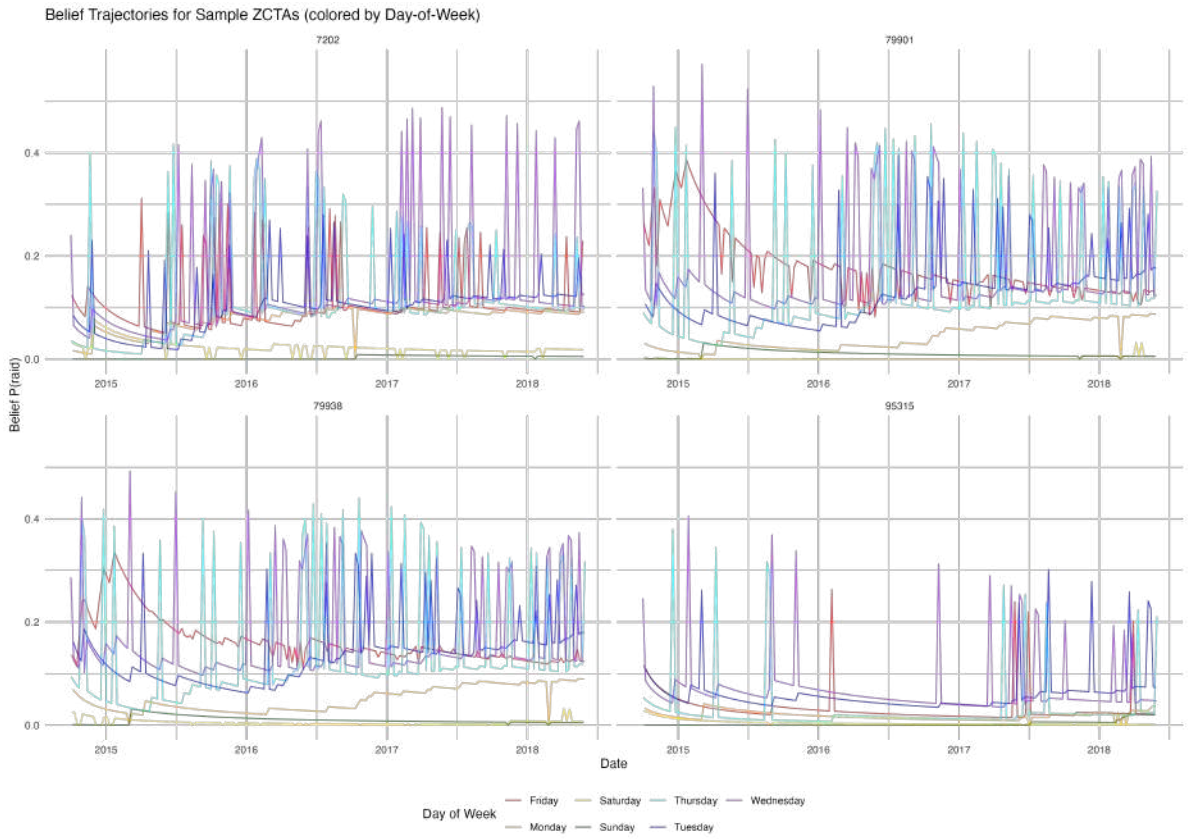
(b) Cumulative forecast error by weekday



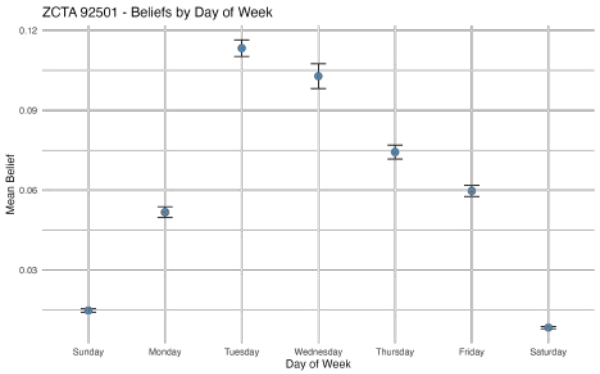
(c) Marginal effect of beliefs over time

Figure A17: Implied Beliefs, Forecast Errors, and Time-Varying Belief Sensitivity

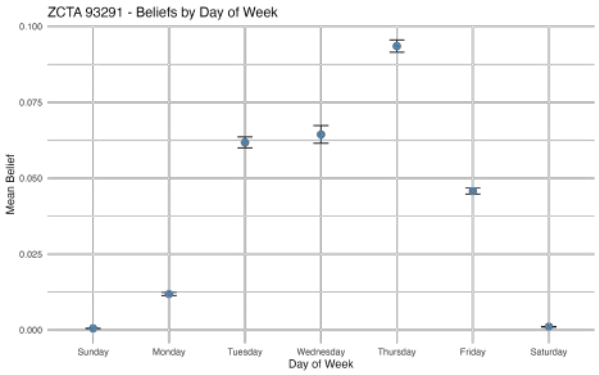
Notes: Panel (a) shows the distribution of implied beliefs $B_{zd(t)}$ by weekday across all ZCTAs in the sample, over time. Panel (b) shows the distribution of implied cumulative forecast errors $\bar{f}_{zt}$ by weekday across all ZCTAs in the sample, over time. Panel (c) displays the distribution of marginal effect of beliefs on economic activity, $\alpha_{11} + \alpha_{12} \cdot \bar{f}_{zt}$, over time, across all ZCTAs in the sample.



(a) Beliefs by weekday over time for a subset of ZCTAs



(b) ZCTA 92501 (Riverside, CA)



(c) ZCTA 93291 (Visalia, CA)

Figure A18: Implied weekday average beliefs for a subset of ZCTAs

Notes: Panel (a) shows the implied belief evolution $B_{zd(t)}$ by weekday over the estimation period for four different ZCTAs. Panels (b) and (c) show the implied mean beliefs by weekday for two example ZCTAs in California.

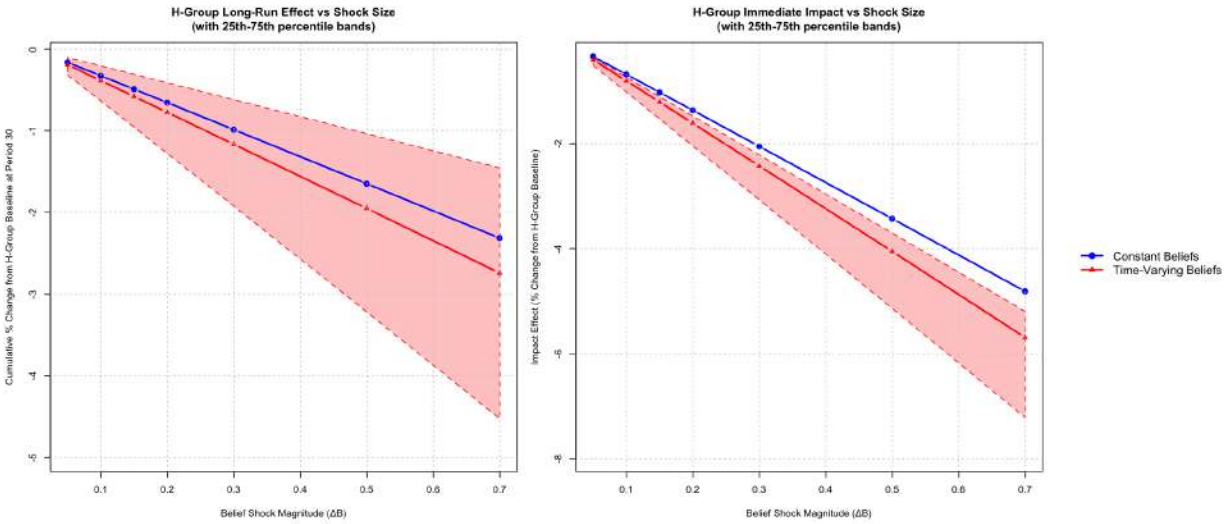
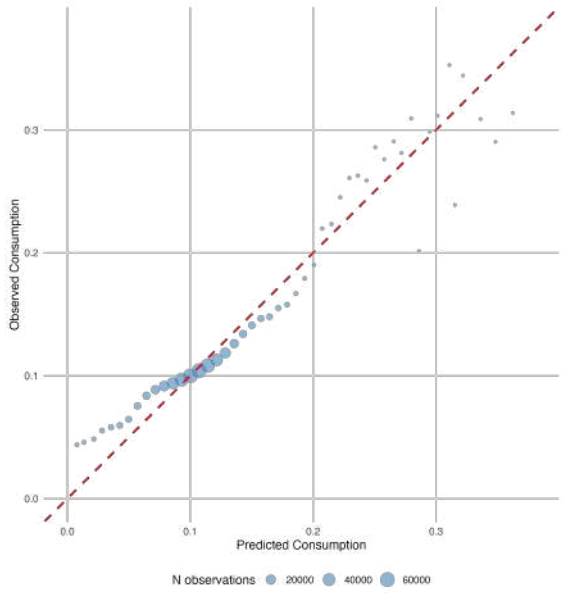


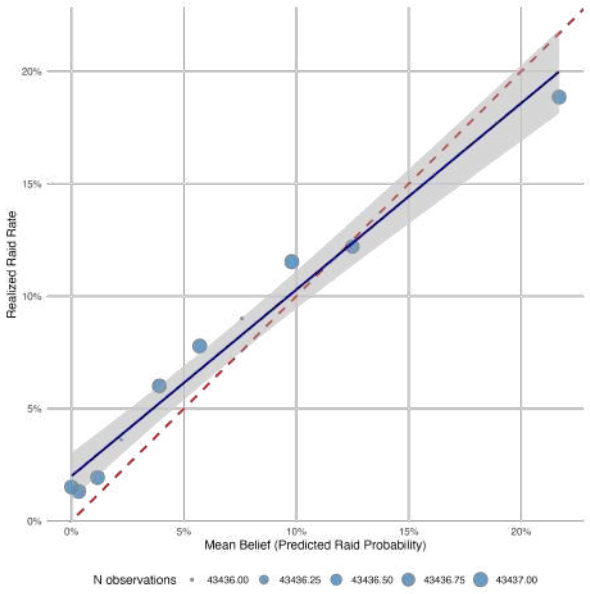
Figure A19: Nonlinearity in Impulse Responses: H-Type Group

Notes: Impulse response effects on economic activity as a percent of the steady state baseline for the Hispanic foreign born group under belief shocks of magnitudes varying between 5 and 70 percentage points. The left-hand-side panel presents the net cumulative effects. The right-hand-side panel presents the effect on impact. The blue lines measure the mean effects under the scenario where beliefs revert to baseline immediately after the shock. The red lines measure the mean effects under the scenario where beliefs are allowed to evolve endogenously after the shock. The pink-shaded areas represent the IQR of the distribution of impulse-response effects across ZCTAs for the scenario where beliefs can evolve endogenously.

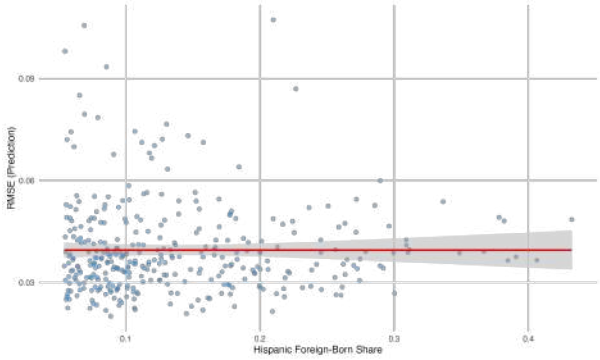
G Model Validation



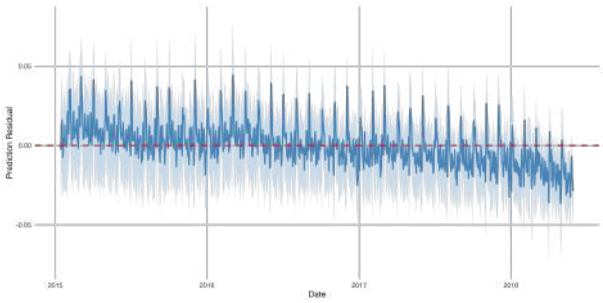
(a) Predicted vs. observed economic activity



(b) Belief calibration



(c) Model fit by Hispanic foreign-born share



(d) Residuals over time

Figure A20: Model Fit Diagnostics

Notes: Panel (a) shows a binned scatter plot of predicted versus observed economic activity (share of accounts active), with 100 equal-sized bins. Panel (b) shows belief calibration: binned mean beliefs plotted against observed raid frequencies. Panel (c) shows RMSE by decile of Hispanic foreign-born share. Panel (d) shows residuals over time, illustrating time-series fit.

Figure A21 presents the results of a placebo test for the structural breaks validation. We generate 233 "fake" structural breaks by randomly reassigning break dates within each ZCTA's sample period, preserving the overall frequency and ZCTA distribution of breaks. The figure overlays the event study coefficients for real breaks (green line) against placebo breaks (grey line). While forecast errors spike sharply after real breaks, the placebo event study is flat—placebo breaks show no effect on forecast errors. This confirms that the

forecast error spikes we observe are not artifacts of our estimation procedure but reflect genuine belief-outcome mismatches caused by actual regime changes.

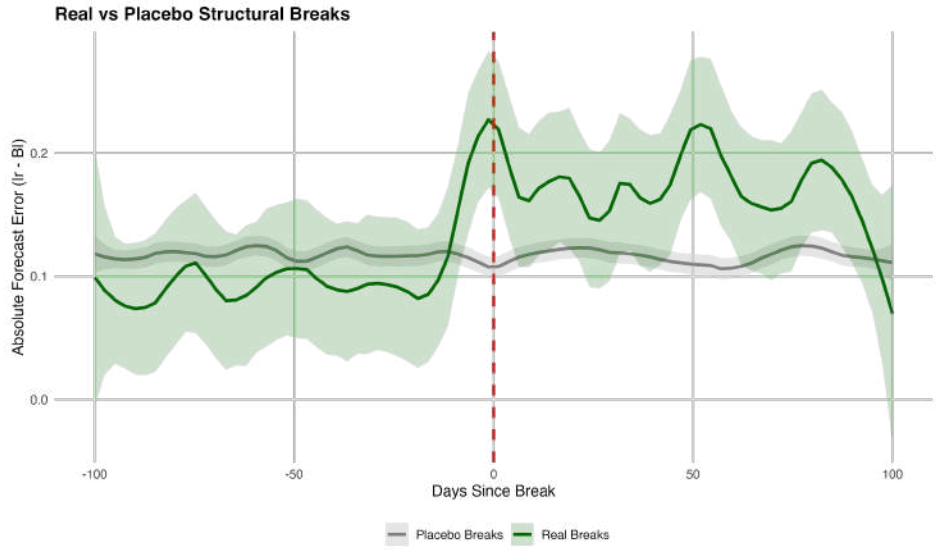


Figure A21: Placebo Test: Real vs. Fake Structural Breaks

Notes: Event study of mean absolute forecast errors around detected structural breaks. The green line shows the conditional mean of the forecast errors from event time based on detected structural breaks following the Bai and Perron (2003) methodology. The grey line shows the conditional mean of the forecast errors from event time based on placebo structural breaks generated by randomly reassigning break dates within each ZCTA. The shaded regions show 95 percent confidence intervals.

	(1) Clustering	(2) Support Ratio	(3) log(SCI)	(4) Cohesion PC1
Cohesion measure	-6.15	-0.70	-0.042	-0.042
	(2.51)	(0.27)	(0.034)	(0.017)
Observations	349	349	349	349
R^2	0.41	0.41	0.40	0.41

Table A3: Structural Model Fit and Social Cohesion

Notes: The dependent variable in all columns is the per-ZCTA RMSE from the structural model's share accounts active residuals. Each column reports the coefficient on a different social cohesion measure: Column 1 uses the clustering coefficient from Chetty et al. (2022), measuring the average across Facebook subscribers in the ZCTA of the share of an individual's Facebook friend pairs who are also friends with each other; column 2 uses the support ratio from Chetty et al. (2022), measuring share of Facebook friendships in the ZCTA that are supported (have at least one mutual friend); column 3 uses the Social Connectedness Index from Bailey et al. (2018), that measures the ratio of Facebook friendships observed between individuals in a given ZCTA as a scaled fraction of all potential friendships between them; column 4 is the first principal component of the three measures. All specifications include log population and log number of raids in the ZCTA as controls. Standard errors in parentheses.

H Counterfactuals

Figure A22 shows the distribution of consumption losses across ZCTAs for each raid increase scenario. The distributions shift left as the enforcement rate increases, with losses becoming more concentrated at higher negative levels. The distribution of losses is bimodal, and spreads out as the increase in enforcement intensity becomes higher.

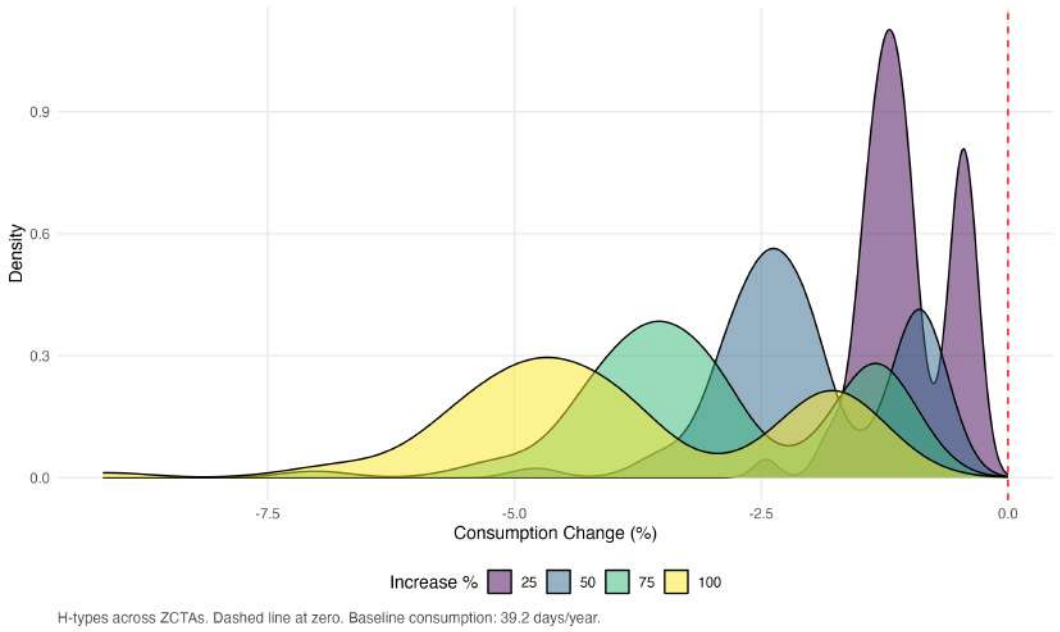


Figure A22: Distribution of Consumption Losses by Raid Increase Level

Notes: Distribution of consumption losses (percent) among the Hispanic foreign born across the 352 ZCTAs under each raid increase scenario. Losses are computed relative to the baseline (0 percent reduction) and averaged over 50 Monte Carlo simulation draws.