The Manipulation of Algorithmic Decisionmaking and the Expansion of Immigration Detention

David Hausman & Michael K.T. Tan[1]

Abstract

In two steps, in 2015 and 2017, ICE changed the algorithmic tool that helps officers decide which noncitizens to detain. It removed the bond and release options, allowing the tool only to recommend detention or to refer cases to a supervisor. Using new data obtained by FOIA request, we document these two changes and show that the second, in 2017, substantially reduced the release rate. We find that both the 2015 and 2017 changes followed periods of decreasing apprehensions, and we hypothesize that ICE may have changed the tool in order to fill empty detention beds.

[1] David Hausman is a fellow at the Stanford Law School Regulation, Evaluation and Governance Lab and a Ph.D. Candidate in the Stanford Department of Political Science. He was an attorney at the ACLU Immigrants' Rights Project from 2016 to 2019, and continues to consult occasionally. He is also a consulting expert on issues related to the RCA tool in New York; his analysis for that case is separate from this article and uses a separately produced dataset. Michael Tan is the Deputy Director of the ACLU Immigrants' Rights Project. This proposal reflects the views of the authors and not the ACLU.

In 2015, Immigration and Customs Enforcement ("ICE") altered its software for recommending detention, bond, or release: it removed the bond recommendation possibility. In 2017, ICE again altered its software: this time it removed the release recommendation possibility. After the 2017 change, the software—intended to help ICE officers decide whether or not to detain noncitizens—could only make two possible recommendations: detain or refer to supervisor. In other words, the algorithm no longer made predictions; it only instructed ICE officers to detain.

This article draws on a new dataset, obtained through a lawsuit under the Freedom of Information Act ("FOIA"), to document this algorithmic manipulation quantitatively and to study its causes and effects. Using new individual-level data on over a million ICE risk classification decisions, we show that ICE altered its risk algorithm in two discrete steps—in February 2015 and June 2017. We also show, with a regression discontinuity design, that the 2017 change decreased noncitizens' chance of release by about half—from around three percent to just over one percent. The 2015 change, by contrast, had little effect, but it followed a steep decline in the bond grant rate in the previous months. Finally, we find that both of these changes followed steep declines in the number of ICE detention decisions, and we hypothesize that ICE changed its algorithm to fill newly empty detention beds. Algorithmic risk prediction tools have received praise for their ability to outperform human predictions and criticism for their capacity to perpetuate bias. Our findings highlight an unrelated risk: that governments might strip those tools of their predictive power and use them instead as blunt instruments to eliminate individualized discretion.

Background and Literature Review

We contribute to the sociolegal literature by asking two causal questions. First, did changes to the software underlying the risk tool cause ICE officers to increase their use of detention? Second, what ICE policy changes led to the changes in the risk tool software? We use a regression discontinuity design to answer the first question, and we find that the 2017 change to the software did cause a (relatively small) change in decisions about whether to detain. We rely on descriptive and qualitative evidence to answer the second question: we find that the software changed more to match ICE policy than vice versa. These results contribute to three sociolegal literatures.

These findings contribute, first, to the quickly growing literature on combined human-algorithmic decisionmaking. A series of studies has demonstrated that algorithms can outperform human judges' predictive judgments about recidivism in the bail context (e.g. Kleinberg et al. 2017). That literature has led to an outpouring of interest in the content of such algorithms for pretrial release and the likelihood that those algorithms perpetuate bias (Corbett-Davies and Goel 2018; Corbett-Davies et al. 2017; Mayson 2017). We are not aware of any work, however, that addresses the question we ask here: whether human adjudicators continue to follow the suggestions of an algorithmic decisionmaking system even after policymakers remove the algorithm's ability to make predictions—here, to recommend release for noncitizens who pose little predicted risk.

Second, these findings add to a rich sociolegal literature examining the causes and consequences of immigration detention (Ryo 2019). Scholars have studied variation in U.S. immigration detention over time and across geographic regions: Ryo and Peacock (2018) find that detention in locations with less access to legal counsel tends to last longer; Eagly and Shafer (2015) document widely varying access to counsel across immigration courts. Our finding that the effect of the 2015 and 2017 changes to the RCA tool varied across field offices is consistent with these results and matches the finding that the use of detention has risen sharply during the Trump administration (Kim & Semet 2019; Ryo & Peacock 2018, p. 9).

In particular, our study contributes to scholarship on how decisions to detain are made. Ryo (2016, 2019a) has found that bond decisions by immigration judges do not take account of many clearly relevant facts, and others have criticized the custody determination process for arbitrarily imposing near-automatic detention (Gilman 2016). The ascendant role that ICE's risk classification tool plays in detention decisions has also been acknowledged and assessed by scholars. Noferi and Koulish (2014) examine an earlier version of the risk classification assessment data that we use in this paper and conclude that ICE engages in significant overdetention, and Koulish & Calvo (2019) use that earlier dataset to examine the contextual determinants of ICE officers' decisions to override RCA recommendations. Koulish (2015) conducts an in-depth assessment of several hundred ICE risk classification decisions and determines that noncitizens who are mandatorily detained are no more likely to pose significant risk than those whom ICE has discretion to release. We build on this past work to document, for the first time, the causes and consequences of the dramatic changes that ICE made to the risk assessment tool in 2015 and 2017.

Finally, our findings also contribute to the literature on the control of discretion. That literature documents many cases in which governmental systems successfully resist new constraints on discretion. For example, the sociolegal literature documents repeated attempts and repeated failures to cabin the exercise of individual discretion in the criminal system (e.g. Miethe

1987; Ulmer et al. 2007; Engen and Steen 2000; but see Yang 2015). In particular, efforts to cabin the discretion of judges through sentencing guidelines appears at least partly to have led prosecutors to widen their exercise of discretion through a process of “hydraulic displacement” (Miethe 1987). We, by contrast, find an instance of the opposite phenomenon: Congress directed ICE officers to exercise their individual discretion, but policies put in place by ICE officials prevented them from doing so. In policy changes in 2015 and 2017, ICE systematically reduced the discretion of individual ICE officers, partly by depriving the RCA software of its predictive capability.

Background: The ICE Risk Classification Assessment

In this section, we briefly survey the history and uses of ICE’s risk assessment tool. We draw on an original source to supplement the existing literature: the deposition of a senior ICE official in a case that we litigated.

The history of ICE’s risk classification tool began in 2009, when Dr. Dora Schiro, Special Advisor to the DHS Secretary, issued a report with a series of recommendations for rationalizing immigration detention (Schriro 2009). Following that report, ICE announced a series of reforms to its detention procedures, including changes to contracting practices, plans to increase the use of alternatives to detention, and the development of a new risk assessment tool to predict whether noncitizens would pose a danger or flight risk if released during their deportation proceedings (ICE 2009). The last reform resulted in the creation of the ICE Risk Classification Assessment (“RCA”), a computerized algorithm that made recommendations about ICE custody decisions by predicting the risk that they would flee or commit crimes.

The algorithm was first put to use in mid-2012. ICE made the use of the software mandatory for all arrested noncitizens who were not both subject to mandatory detention[2] and likely to be deported within five days (OIG 2015, 11).[3] Because collecting the information required for the software was time-consuming, however, ICE allowed its officers to wait up to five days after an arrest to make the determination (OIG 2015, 11). The risk classification assessment tool made two types of recommendations. (In this article, we examine only the first of these recommendations.) First, it recommended whether ICE should detain the noncitizen, release the noncitizen on bond, release the noncitizen on community supervision (without a bond), or whether an ICE supervisor should make the decision in the first instance. Second, the software evaluated what level of custody risk each person posed. The RCA procedure required two levels of review of the software's recommendations: first, an ICE officer could agree or disagree with the tool's recommendations, and second, an ICE supervisor could agree or disagree with the ICE officer's recommendation.

[2] For certain categories of noncitizens, ICE does not have a choice about detention. Congress has decided that all noncitizens convicted of certain crimes must be detained pending their proceedings (8 U.S.C. 1226(c)).
[3] Whether ICE field offices in fact treat the use of the tool as mandatory is an open question. See n.6 below.

The risk assessment procedure fit into the existing statutory framework for ICE decisions about whether to detain noncitizens. Section 236(a) of the Immigration and Nationality Act, 8 U.S.C. § 1226(a), provides that, while noncitizens who are not subject to mandatory detention are contesting their deportation, the government may decide whether to detain or release them. The implementing regulations have required this decision to take place in two steps, one in the Department of Homeland Security (where the RCA tool is used) and the other in the Department of Justice (where an Immigration Judge makes a decision). In the first step, an ICE officer makes a decision about detention within 48 hours of arrest. 8 C.F.R. § 287.3(d). That decision is discretionary: the officer "may, in the officer's discretion, release an alien . . . provided that the alien must demonstrate to the satisfaction of the officer that such release would not pose a danger to property or persons, and that the alien is likely to appear for any future proceeding." 8 C.F.R. § 1236.1(c)(8). If the ICE officer decides to order that the noncitizen remain detained, the noncitizen can appeal that decision to an Immigration Judge in the Department of Justice (and from there, to the Board of Immigration Appeals). *See* 8 C.F.R. §§ 1003.19(a), 1236.1(d).

In the first year and a half of the RCA tool's use, ICE officers overrode its recommendations more than twenty percent of time (OIG 2015, 11-12). A scathing report from the Office of the Inspector General concluded that the software tool was "time consuming, resource intensive, and not effective in determining which aliens to release or under what conditions" (OIG 2015, 11). The Inspector General emphasized the high override rate, the many questions asked by the software, and the fact that ICE had never conducted any testing of whether the tool correctly predicted risk (OIG 2015, 12). Finally, the Inspector General also noticed that the tool made no recommendation at all in nearly twenty percent of cases (OIG 2015, 11).[4]

From 2013 through 2015, the override rate decreased to under 10% (Koulish and Calvo 2019, Figure 2). Koulish and Calvo, as well as Schriro (2019), suggest that this decreasing dissent rate resulted from ICE changes to the software resulting in fewer release or bond recommendations. Schriro (2019, 41-42) concludes that "[b]y eliminating factors key to assessing risk and replacing them with policy-based measures, the RCA's algorithm lost the ability to measure true risk," and that "the risk recommendation was reverse engineered to reflect enforcement personnel's preference for detention."

As we show in detail below, this process continued in 2015 with two significant, discrete changes to the RCA software. In February 2015, the RCA software ceased to have the capability to recommend that a noncitizen be released on bond, and in June 2017, the RCA software ceased to have the capability to recommend that a noncitizen be released on community supervision (Reuters 2017). As a result, by 2017, the RCA tool could make only two recommendations: that a noncitizen be released, or that a supervisor make the custody decision in the first instance.

We rely on an additional, original source of qualitative evidence to confirm these changes and the fact that they were made not for predictive reasons but instead in response to ICE decisions to detain more people. In a deposition in *Hernandez v. Barr*, No. 16-620 (Central District of California), the unit chief of ICE's Information Technology Management Unit confirmed that changes to the RCA software removed the bond recommendation (Wilson 2019,

[4] We replicate this problem, finding that, over the full period of our dataset, the RCA tool failed to make any risk prediction in 384,519 of 1,348,363 cases. We believe that these missing recommendations reflect expedited removal cases. The missing predictions begin in August 2013, and ICE noted in its response to the Office of the Inspector General that it, in that month, "streamlined the RCA by generating an automatic detain decision in expedited removal cases, allowing field offices to skip the submission/approval steps otherwise required" (OIG 2015, 13). We have verified that the proportion of missing predictions did not change sharply at the same time as the two software changes that we study below.

63-64). Mr. Wilson also said that the second change to the RCA tool was a result of President Trump's changes to interior enforcement priorities—in other words, that the removal of the release recommendation reflected a substantive policy change (Wilson 2019, 86-87). Finally, Mr. Wilson admitted that the goal of the tool was not to improve risk predictions but instead to approximate the predictions that officers would make in the absence of the tool: "The goal would be a lower override rate because that means that you've built a tool that is mimicking the decisions that the officer would normally make. So, for instance, when the memo came out and officers were directed to take certain actions, that memo came out, they started taking those actions, then the RCA was changed -- RCA was changed to mimic those actions they were already taking."

In sum, the qualitative evidence indicates that ICE used the RCA software not to improve risk predictions but instead to implement its detention priorities. We offer evidence below that bed space, in addition to changed enforcement priorities, helped drive these two changes. We also test whether the changes to the software had an effect on human decisionmakers.

Data and Descriptive Statistics

The data come from a Freedom of Information Act Request from the American Civil Liberties Union to Immigration and Customs Enforcement in 2019 (2019-ICFO-10844). ICE produced the data after the ACLU filed a complaint in the District Court for the Southern District of New York (Case No. __). The case remains in litigation; ICE has not produced all of the data that the ACLU requested. As part of its initial production, however, ICE did produce individual-level data on all Risk Classification Assessment decisions from mid-July 2012 to the end of August 2019.[5]

We present basic descriptive statistics on detain/release recommendations and decisions (*N* = 1,348,363).[6] The RCA tool makes one of four recommendations: detain, release, bond, or supervisor to determine. The RCA final decision includes only three options—to detain, release, or grant bond—since the supervisor has by that point made his or her determination. Table 1 shows numbers of decisions and mean detention rates by year. Table 2 shows similar statistics by ICE field office.

Table 1: Descriptive Statistics by Year

Year	Mean Detention Decision Rate	Mean Detention Recommendation Rate	N
2012 (after July)	75.01%	46.38%	11862
2013	78.81%	53.42%	211022
2014	84.23%	71.70%	183498
2015	93.68%	81.94%	155757
2016	94.18%	84.04%	208378
2017	94.37%	85.74%	198325
2018	93.77%	92.28%	227930
2019 (through August)	96.42%	95.37%	151591
Total	90.40%	79.91%	1348363

[5] That data includes more than one type of risk classification decision. As explained *supra*, ICE uses the RCA tool not only to determine whether to detain or release individuals it has apprehended, but also to determine what risk level they might pose in custody.

[6] Because the RCA tool is not required where the officer believes that the noncitizen is subject to mandatory detention and will be deported within five days; ICE may in these cases apprehend a person and make a custody decision without using the RCA tool at all. That means that the RCA data does not offer a comprehensive measure of ICE's detention decisionmaking. Moreover, Mr. Wilson said in his deposition that field offices have the discretion not to use the RCA tool at all (Wilson 2019, 26). We are unsure how often that occurs. In any event, that possibility does not undermine the research strategy below, which relies on that dataset to make inferences about whether changes to the RCA tool caused sharp changes in ICE officers' decisionmaking *when using the tool*.

The yearly data in Table 1 show increasing detention over time. While both the actual detention rate and the recommended detention rate have increased, the gap between the two has decreased over time. In 2012, while 75% of assessments resulted in detention, the RCA tool only recommended detention 46% of the time. Since then the RCA tool has recommended detention more and more often, eventually recommending detention 95% of the time in 2019. While the regression discontinuity results focus on the changes occurring in 2015 and 2017, Table 1 shows the long-term trend.

Table 2: Descriptive Statistics by Field Office

Field Office	Mean Detention Decision Rate	Mean Detention Recommendation Rate	N
Atlanta	90.96%	74.97%	60090
Baltimore	72.70%	68.63%	10871
Boston	87.11%	70.64%	17466
Buffalo	78.96%	71.65%	14662
Chicago	81.86%	74.11%	50519
Unknown	78.57%	78.57%	14
Dallas	82.97%	82.31%	44835
Denver	82.18%	72.00%	13541
Detroit	92.11%	70.26%	22405
El Paso	96.16%	89.51%	46530
Houston	98.30%	83.15%	84406
Los Angeles	86.25%	65.44%	37253
Miami	91.46%	76.73%	55625
New Jersey	79.85%	64.96%	21370
New Orleans	86.62%	75.90%	39429
New York	80.42%	68.15%	20113
Philadelphia	86.36%	71.14%	26516
Phoenix	94.56%	81.47%	135351
Seattle	93.08%	69.62%	21468
San Francisco	78.37%	69.75%	24979

Salt Lake City	92.03%	83.13%	28659
San Antonio	96.59%	88.45%	413766
San Diego	77.03%	72.83%	109570
Saint Paul	79.98%	75.10%	23936
Washington	84.64%	67.26%	24989
Total	90.40%	79.91%	1348363

In Table 2, actual detention decision rates and detention recommendation rates were high in every field office, but there is significant variation across field offices nonetheless. Notably, the gap between recommended detention and actual detention also varies: for instance, while actual detention mirrors recommended detention in Dallas, actual detention exceeds recommended detention in Houston by over 15 percentage points.

Finally, we also look at simple trends over time. ICE risk assessment decisions—decisions about whether to take someone into continued custody after an initial arrest—changed sharply around two discrete dates: February 16, 2015 and June 5, 2017. On February 16, 2015, the RCA software changed and stopped recommending release on bond; on June 5, 2017, the RCA software again changed, this time to stop recommending release of any kind. ICE officers' human decisions changed around the time of the changes to the software. Figure 1 shows these trends.

Figure 1: Trends Over Time In ICE Risk Classification Decisions

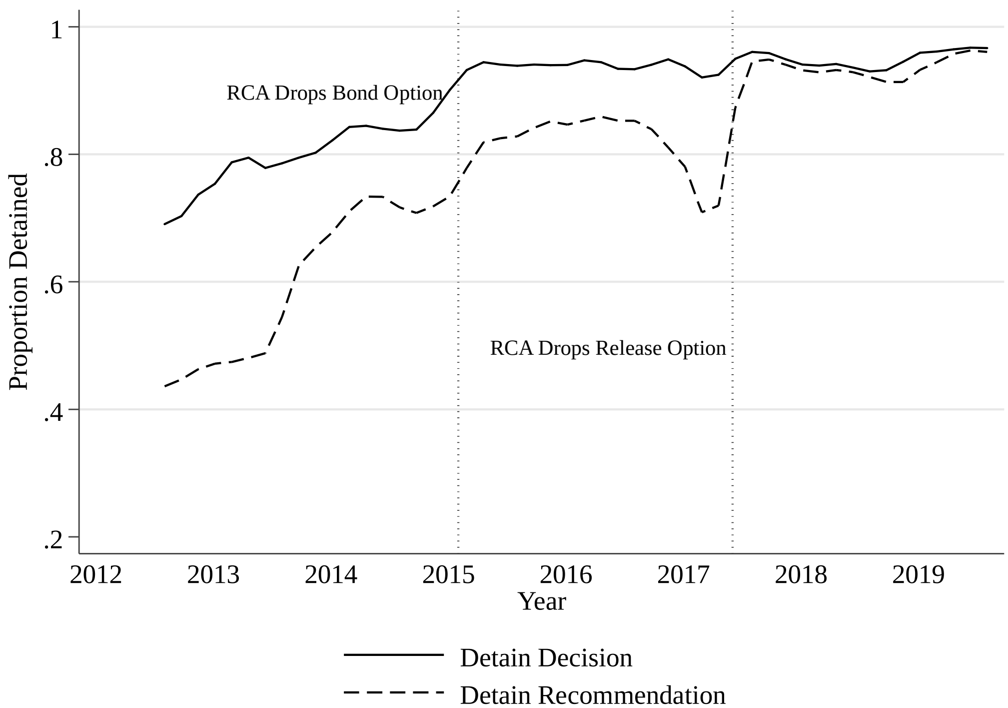

The dotted line shows the proportion of cases in which ICE's algorithm recommended continued detention, and the solid line shows the proportion of cases in which an ICE officer in fact decided to subject the person to continued detention. When the software version first changed, in 2015, the percentage of cases in which the software recommended continued detention jumped from about 60% to about 70%, and the percentage of cases in which ICE officers actually ordered continued detention jumped from about 80% to 90%. When the software changed again, in 2017, the percentage of cases in which the software recommended detention jumped all the way up to 90%, and the percentage of cases in which officers ordered detention jumped, too—but only by a few percentage points, maybe because the rate was already so near 100% that it could barely rise further.

Methods and Results

On a closer look, the story is more complicated. In order to test more formally whether changes to the Risk Classification Assessment tool had an effect on (1) ICE detention decisions and (2) the number of bond hearings before immigration judges, we employ a regression discontinuity ("RD") design with time as the running variable. This is a good context for a time RD design for at least two reasons. First, we have high-frequency data: beginning in 2015, a median day included over 400 risk classification decisions. Second, we know exactly when the changes to the risk classification tool took effect because the version of the tool that ICE employed is recorded in our dataset. In addition, we observe not only risk classification decisions, but also risk classification *recommendations*, and bond and release recommendations each stopped occurring on a discrete day.

In order for the RD design to identify a causal effect, we must assume that no other change occurred at the same time as the changes to the RCA tool. That assumption is plausible here: we have no reason to believe that any other policy changes occurred on February 16, 2015 or June 5, 2017.

Removing the Bond Recommendation

To evaluate whether the removal of the bond recommendation had an effect, we begin by checking that the change in the tool's recommendations was sudden and large; Figure 2 performs this check, showing that the rate at which the tool recommended bond suddenly dropped from around 7% on February 15 to 0% on February 16, 2015. Each dot represents a single day; the dotted lines and gray bands show the smoothed trend, along with 95% confidence intervals, before and after the change in the tool.

Figure 2: RCA Bond Recommendation Rate Before and After Change to RCA Tool

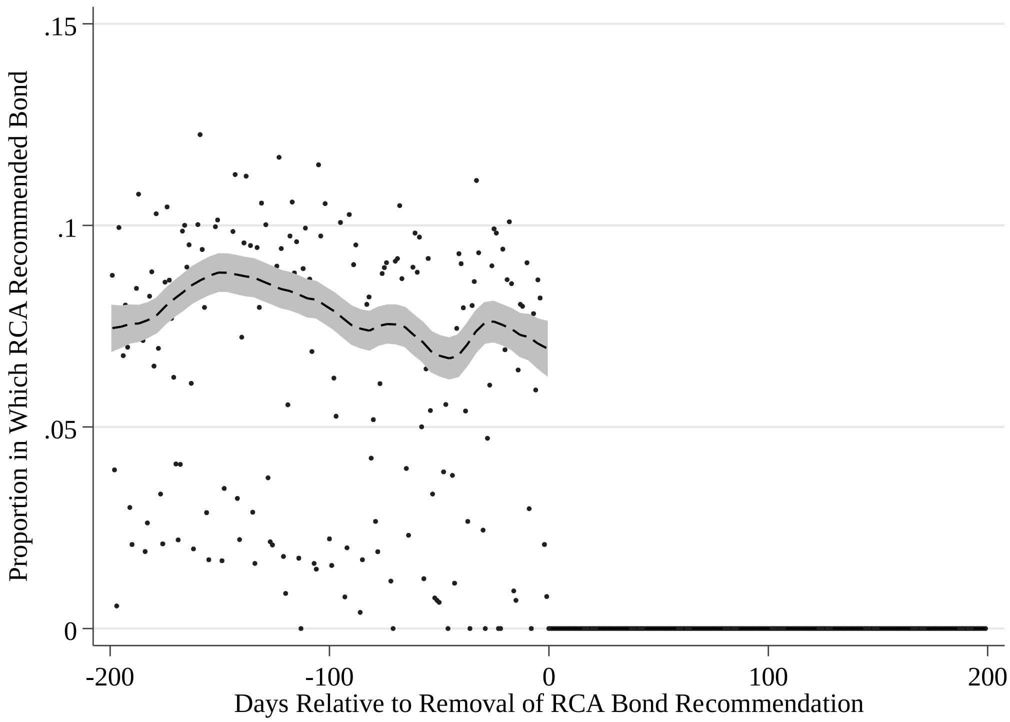

After verifying this sudden drop, we investigate whether actual decisions (not just recommendations) to grant bond changed sharply at the same time—in other words, whether ICE officers followed the new recommendations. If, for instance, ICE officers had universally deferred to the new recommendation, bond grant rates would have fallen to zero. Figure 3 below plots the daily rate at which ICE granted bond against the number of days before and after February 16—the day on which the RCA tool dropped its bond recommendation. Once again, each dot is a day, and the lines show trends and confidence intervals.

Figure 3: Bond Grant Rate Before and After Change To RCA Tool

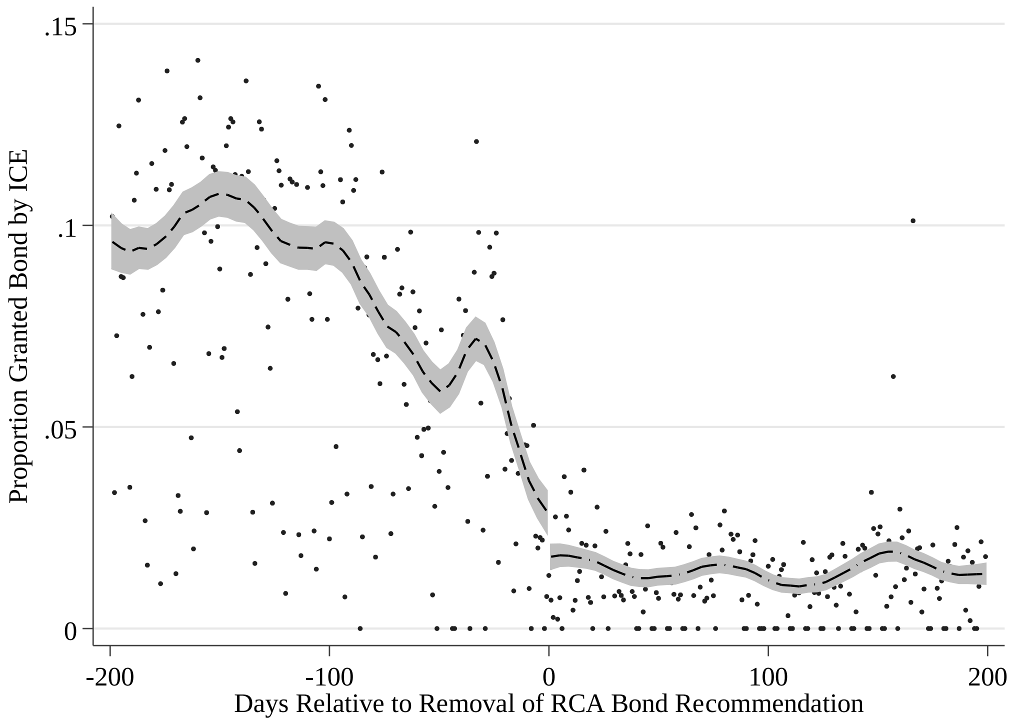

Figure 3 suggests a more complex story than Figure 1 above. In the context of several years of data in Figure 1, the change in release rates appeared sharp and sudden, but when we focus on the six months or so before and after the change, it is clear that ICE officers became less likely to grant bond in the month or two before the change in the RCA tool's recommendations. On the other hand, in addition to this trend downwards, there appears to have been a smaller sharp change on the day of the change in recommendations.

In other words, this figure suggests to us that two things happened around this time. First, ICE gradually changed its policy to reduce how often it granted bond, and the rate of bond grants dropped from around 10% to around 3% over the course of about three months. Second, it removed the bond recommendation from the RCA tool, which looks like it coincided with a sharp drop in the bond grant rate of maybe 1 percentage point on the day of that change.

Although that drop was small in absolute terms, it was large in relative terms—since the rate before than second drop was only 3%, a 1 percentage point drop is a 33% decrease.

The first of these two effects is consistent with what we learned from deposing the ICE official responsible for this program: that ICE made changes to RCA tool over time to more closely match its officers' decisionmaking. (See discussion *supra*.) But the second effect, if it exists, would suggest that these changes to the tool tell only part of the story: the tool may also have affected ICE officers' behavior.

In order to test more rigorously for the second effect—the sharp decrease in bond issuance when the RCA tool was introduced—we turn now to regression discontinuity results. To perform this test, we estimate a linear probability model of the form:

$$Y_{ij} = \beta_0 + \beta_1 d_{ij} + \beta_1 B_{ij} + e_{ij}$$

where Y_{ij} is the outcome (1 if bond is granted, 0 otherwise) for case i in field office j, d_{ij} is the number of days before or after February 16, 2015, and B_{ij} is an indicator variable for whether the new version of the RCA tool had taken effect. Table 3 shows results. Unlike the results from the nonparametric visualization, the results here suggest that the sharp change to the new RCA version had little effect. Table 3 shows the results using bandwidths from 2 days to 30 days; none suggests that the sharp imposition of the new version of the RCA tool had a statistically significant effect.

Table 3: Regression Discontinuity Results: Overall Effect of Removal of Bond Option

Dependent Variable = 1 if Bond is Granted

	(1) 2 Days	(2) 5 Days	(3) 10 Days	(4) 15 Days	(5) 20 Days	(6) 25 Days	(7) 30 days
Days Before/After Change	-0.0072 (0.0055)	-0.0014 (0.0023)	0.00073 (0.00085)	-0.00089* (0.00041)	-0.00088** (0.00030)	-0.0015*** (0.00022)	-0.0015*** (0.00019)
Removal of Bond Recommendation	0.018 (0.017)	-0.0016 (0.015)	-0.018 (0.011)	-0.0046 (0.0080)	-0.0045 (0.0074)	0.0032 (0.0068)	0.0024 (0.0064)
N	737	2620	4926	8475	11286	14455	16362

Standard errors in parentheses
Models include field office fixed effects
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In sum, Table 3 suggests that there is little evidence that the removal of the bond option from the RCA tool had any nationwide effect.

The lack of evidence of a nationwide effect does not mean that the change had no effect in some field offices, however. Since we know that policy often differs systematically by field office, it makes sense to examine field-office-level effects as well. We therefore examine the same pattern by field office. The Appendix shows the graphs for several large field offices; here, we show two extreme examples. Some field offices show little change at the time that the RCA tool was introduced, while others show significant jumps. In El Paso, for example, the trend in bond decision rates is nearly smooth across the time of the change (see Figure 4). In New York, by contrast (Figure 5), the introduction of the tool seems to mark a sudden shift. Most of the field offices—as the figures in the Appendix show—are somewhere in between these extremes: they show a downward slope before the change to the RCA tool but also show a sudden drop after that change.

Figure 4: El Paso Bond Decisions Before and After 2015 Change to RCA Tool

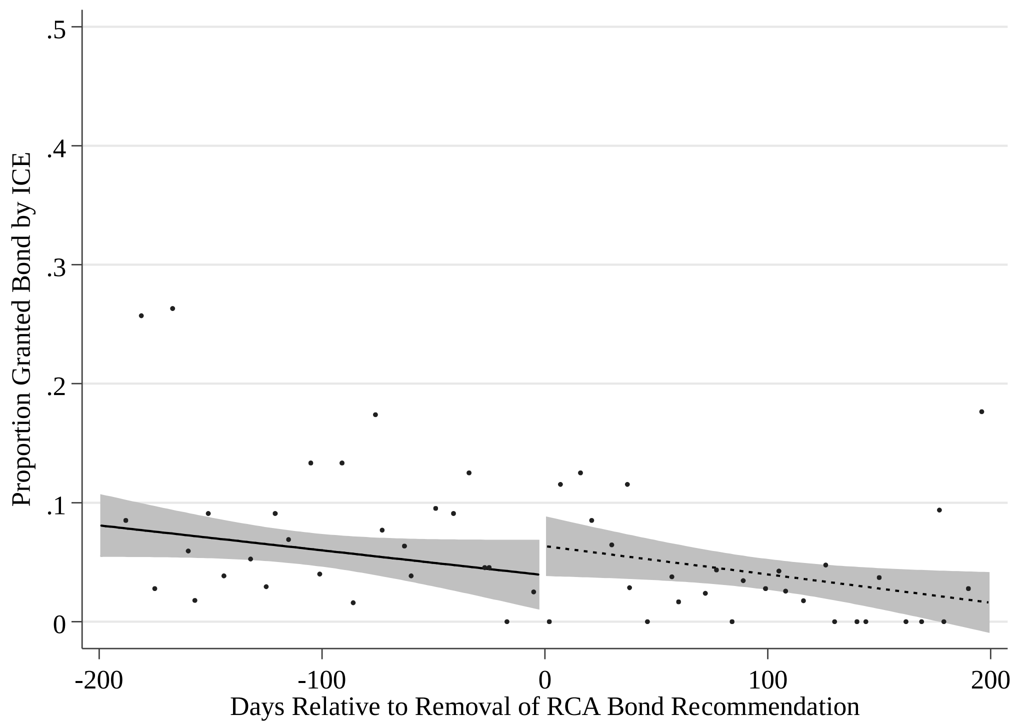

Figure 5: New York Bond Decisions Before and After 2015 Change to RCA Tool

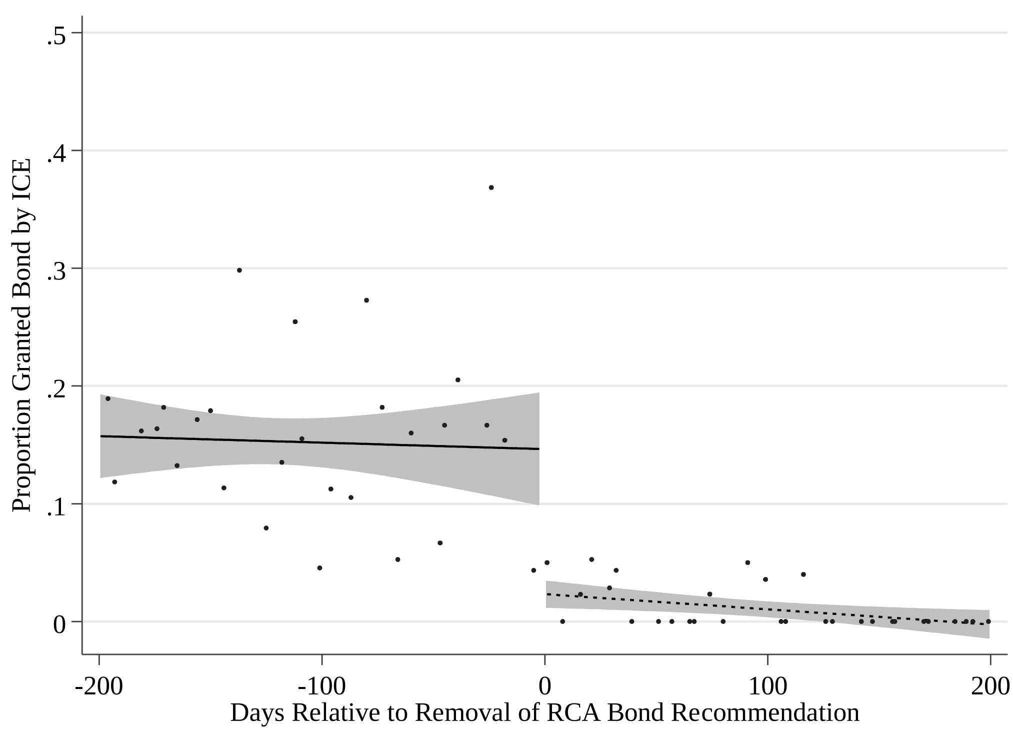

Considering this range of field-office-level results, we next test whether the change in the RCA tool had a larger effect in field offices that granted bond at higher than average rates before change. The logic here is simple: field offices that rarely granted bond even before the removal of the bond recommendation were unlikely to experience large changes, since they did not have far to fall. To test this hypothesis, we estimate the same regression discontinuity models for a subset of field offices with relatively high bond grant rates in the month before the switch to the new tool. To isolate these field offices, we look specifically at how often field offices granted bond *conditional on* a bond recommendation from the RCA tool. We limit our focus to field offices that granted bond at least 25% of the time in this situation (in the 30 days leading up to February 16, 2015).

Table 4: Effect of RCA Tool Change for Field Offices With 25% or Higher Bond Grant Rates (Given Bond Recommendation) in Month Before Change to RCA Tool

Dependent Variable = 1 if Bond is Granted

	(1) 2 Days	(2) 5 Days	(3) 10 Days	(4) 15 Days	(5) 20 Days	(6) 25 Days	(7) 30 days
Days Before/After Change	-0.014^{*} (0.0067)	-0.0024 (0.0032)	0.00028 (0.0012)	-0.0018^{**} (0.00058)	-0.0015^{***} (0.00042)	-0.0020^{***} (0.00031)	-0.0021^{***} (0.00027)
Removal of Bond Recommend ation	0.044^{*} (0.020)	0.0015 (0.021)	-0.017 (0.015)	-0.00019 (0.011)	-0.0033 (0.010)	0.0020 (0.0097)	0.0024 (0.0092)
N	480	1638	3077	5252	7047	8991	10156

Standard errors in parentheses
Models include field office fixed effects
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Again, we find inconsistent effects; in fact, for the smallest bandwidth, we find the opposite of the expected effect—field offices with high grant rates are more likely to grant bond just after the change to the tool. We do not put much weight on a result from just one specification, and we therefore view these results, together, as suggesting that the 2015 change had little effect.

Unfortunately, we are unable to go into more detail still and examine these effects field office by field office, since we lack enough observations within each field office in the smaller bandwidths. Still, our visualizations for some of the field offices are consistent with a larger effect than the nationwide effect. In the New York field office, for example (Figure 5 above), the visual evidence suggests that the change to the tool had around a 10 percentage point effect.

Taking the evidence together, we find little evidence that the February 2015 change to the RCA tool had a small effect on ICE officials' custody decisions. We're confident, however, that the change to the RCA tool in that month was driven by a swift *preceding* change in ICE bond practices. This finding follows from the descriptive and qualitative evidence above. But what drove the swift change that the RCA tool adapted to? Here, we are forced to speculate.

The Inspector General suggested in its 2015 report that some ICE detention decisions are motivated more by whether ICE has sufficient bed space than by whether a given individual poses a flight risk or a danger (OIG 2015, 13). When ICE has more bed space available, it can detain more people, and therefore has the freedom to detain a larger proportion of the people it arrests. Available bed space could explain the strong trend in detention decisions before the changes to the RCA tool. Figure 6 shows how many detention decisions ICE made each week in the time before and after the change to the RCA tool. In the few months leading up to the change, the number of weekly decisions declined from a high about 3,500 to a low below 2,500—a more than 30% decline. That decline almost certainly meant that more beds were available in ICE detention facilities; that bed space might have led ICE officials to take steps to detain more of the people they arrested. As ICE had more and more bed space toward the end of 2014 and the beginning of 2015, it became less and less likely to release people on bond.

Figure 6: Number of Decisions Before and After 2015 Change to RCA Tool

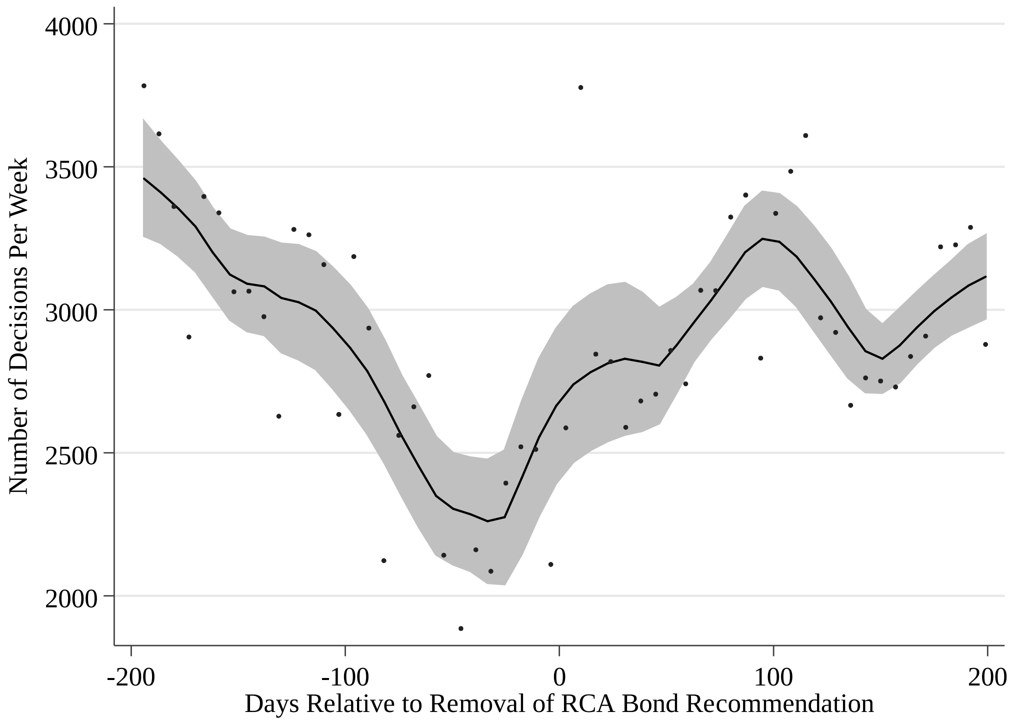

Removal of the Release Recommendation from the RCA Tool

The removal of the bond recommendation was the first of the two major changes to the tool that we study here. The second change occurred on June 5, 2017 when ICE removed the "release" recommendation from its detain/release RCA tool. Looking back at Figure 1 above, there was a large spike in June 2017 in the rate at which the RCA tool recommends detention and an accompanying, but a much smaller spike in the rate at which ICE officers actually ordered detention. (Since the detention rate was already above 90%, it did not have that much further up to go.)

To evaluate the effects of this second change to the RCA tool, we use the same methods we used for the first. We start with a visual inspection of trends in the RCA tool's release recommendations (Figure 7) and ICE officials' release decisions (Figure 8). In Figure 7, the effect of the removal of the release option is obvious: release recommendations drop out of the data almost entirely. And in Figure 8, unlike in the case of the 2015 change, the trend in releases is relatively flat before the change, and makes a sudden fall downward at the time of the change. Although this decline is small in absolute terms, it is large in relative terms—this two to three percentage point drop represents a relative decrease of about half.

Figure 7: Release Recommendation Rate Before and After 2017 Change to RCA Tool

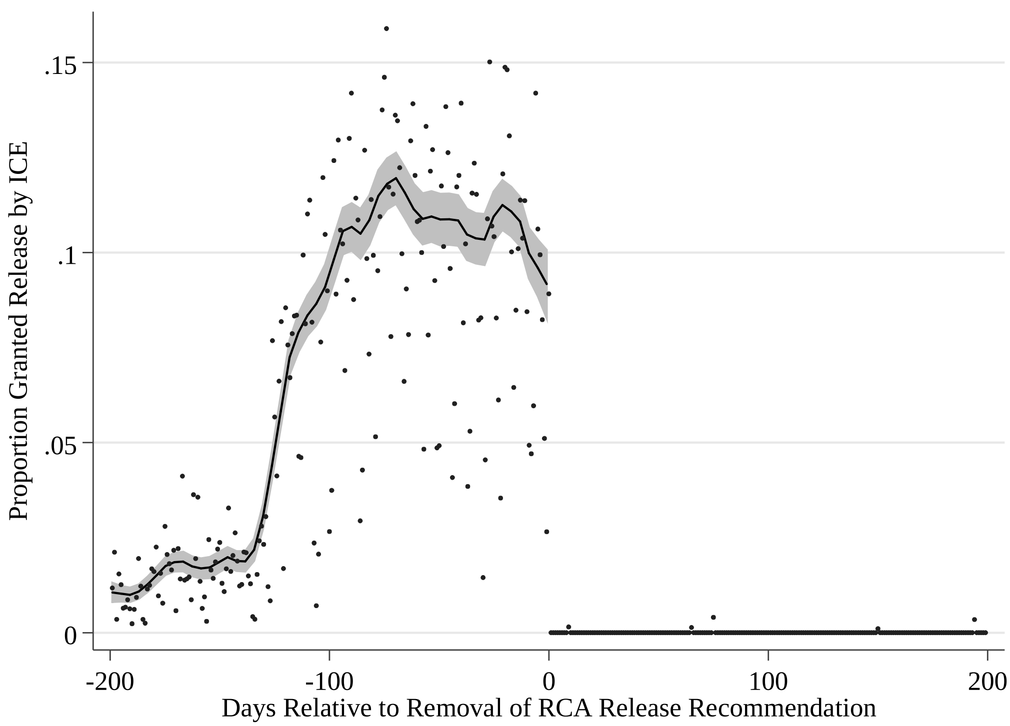

Figure 8: Release Decision Rate Before and After 2017 Change to RCA Tool

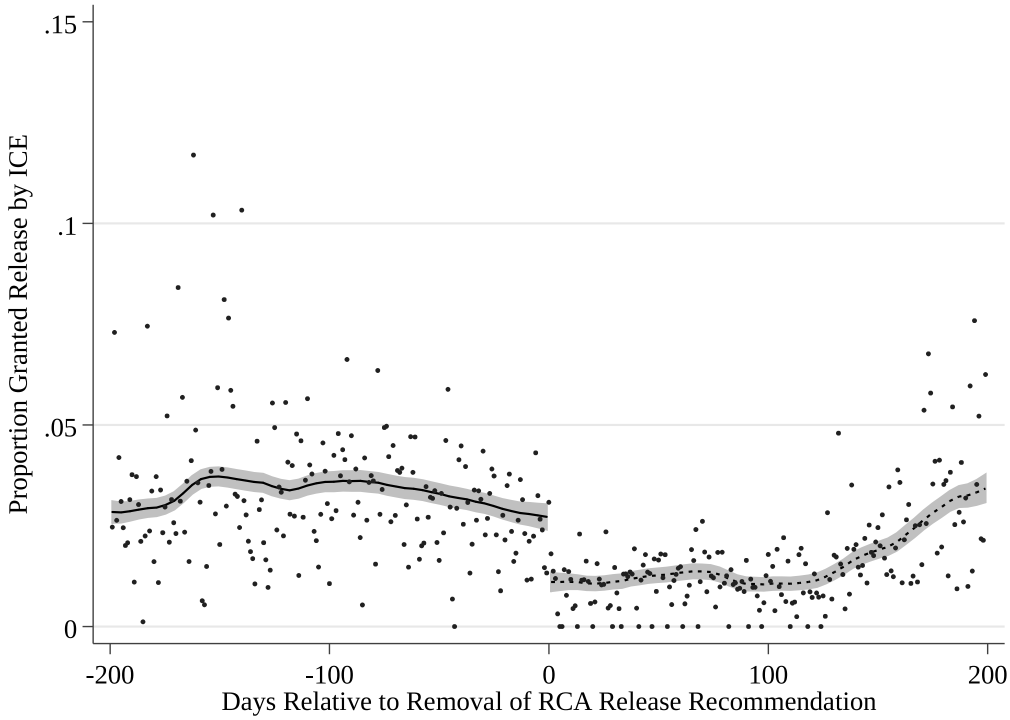

Again, to test more formally whether the change to the RCA tool caused this fall, we estimate a linear probability model of the form:

$$Y_{ij} = \beta_0 + \beta_1 d_{ij} + \beta_1 B_{ij} + e_{ij}$$

where Y_{ij} is the outcome (1 if release is granted, 0 otherwise) for case i in field office j, d_{ij} is the number of days before or after June 5, 2017,[7] and B_{ij} is an indicator variable for whether the new version of the RCA tool had taken effect. Table 5 shows the results. Here, likely because there was not a strong trend in the release rate before the change to the RCA tool, the results are unambiguous. All but one of the bandwidths suggest that the change to the RCA tool caused a 2-3 percentage point reduction in releases. Again, this is a large decline in relative terms, given that the release rate before the change hovered under 5%.

[7] Note that on June 5 itself, some determinations used version 6.3 while others used version 6.4, and a tiny handful of determinations used version 6.4 before June 5.

Table 5: Regression Discontinuity Estimates of the Effect of the 2017 Change to the RCA Tool

Dependent Variable = 1 if Release is Granted

	(1) 2 Days	(2) 5 Days	(3) 10 Days	(4) 15 Days	(5) 20 Days	(6) 25 Days	(7) 30 days
Days Before/After Change	0.0018 (0.0050)	-0.0016 (0.0015)	0.00042 (0.00052)	0.00031 (0.00030)	0.00026 (0.00020)	0.000074 (0.00014)	-0.000018 (0.00012)
Removal of Release Recommendation	-0.029^{*} (0.012)	-0.014 (0.0091)	-0.027^{***} (0.0062)	-0.027^{***} (0.0053)	-0.026^{***} (0.0048)	-0.023^{***} (0.0041)	-0.021^{***} (0.0040)
N	1942	4802	8328	12270	16067	20529	23327

Standard errors in parentheses
Models include field office fixed effects
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Moreover, this effect once again varied strongly across field offices, with some barely affected and others experiencing larger than average effects. Figures 9 and 10 show examples; San Antonio's release rate, which was already nearly indistinguishable from zero, barely changed, whereas the release rate in Washington dropped by more than five percentage points.[8]

[8] Again, in the (online) Appendix we show these figures for every field office. (That Appendix is included in this document now for convenience.)

Figure 9: San Antonio Release Decisions Before and After 2017 Change to RCA Tool

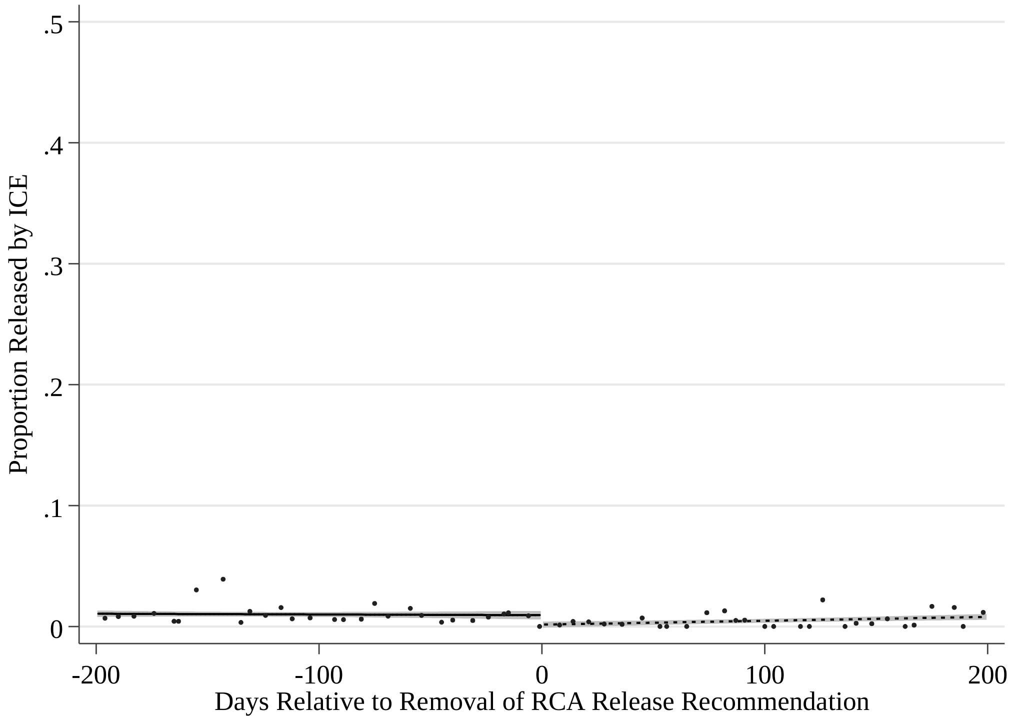

Figure 10: Washington Release Decisions Before and After 2017 Change to RCA Tool

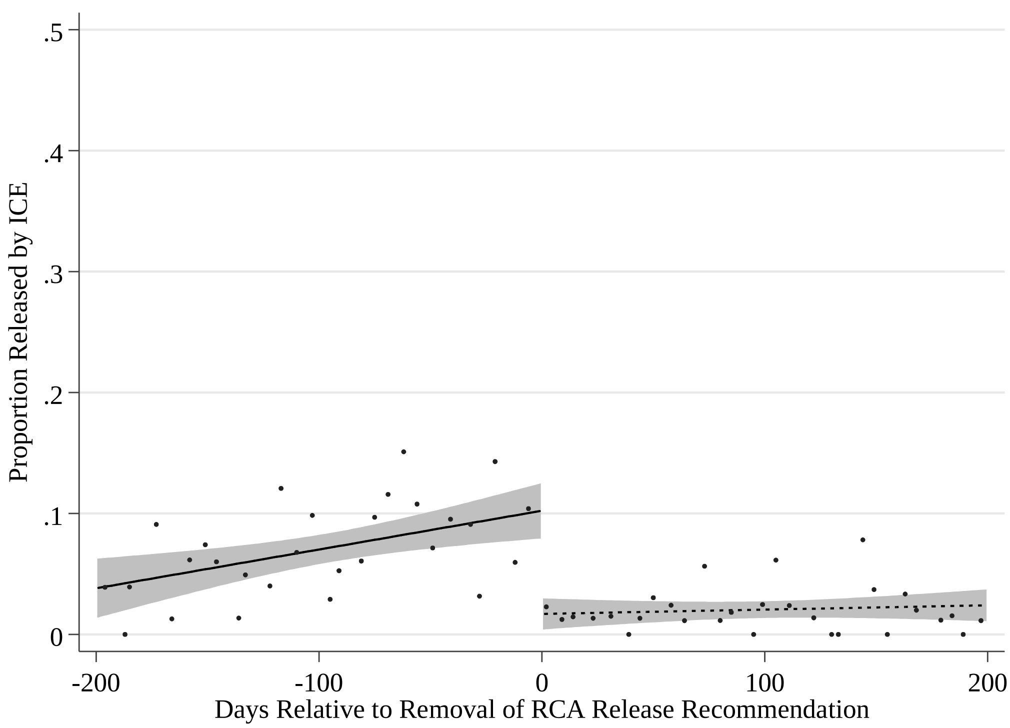

What caused this 2017 change in policy? In our deposition, the ICE official said that this change took place as a result of President Trump's new enforcement priorities, and we see no reason to doubt that statement. As in the case of the 2015 change, however, we graph the number of decisions per week leading up to the change (Figure 11). Once again, we find that the change followed a steep decline in the number of decisions reached by ICE—and therefore likely an increase in available detention bed space.

Figure 11: Number of Decisions Before and After 2017 Change to RCA Tool

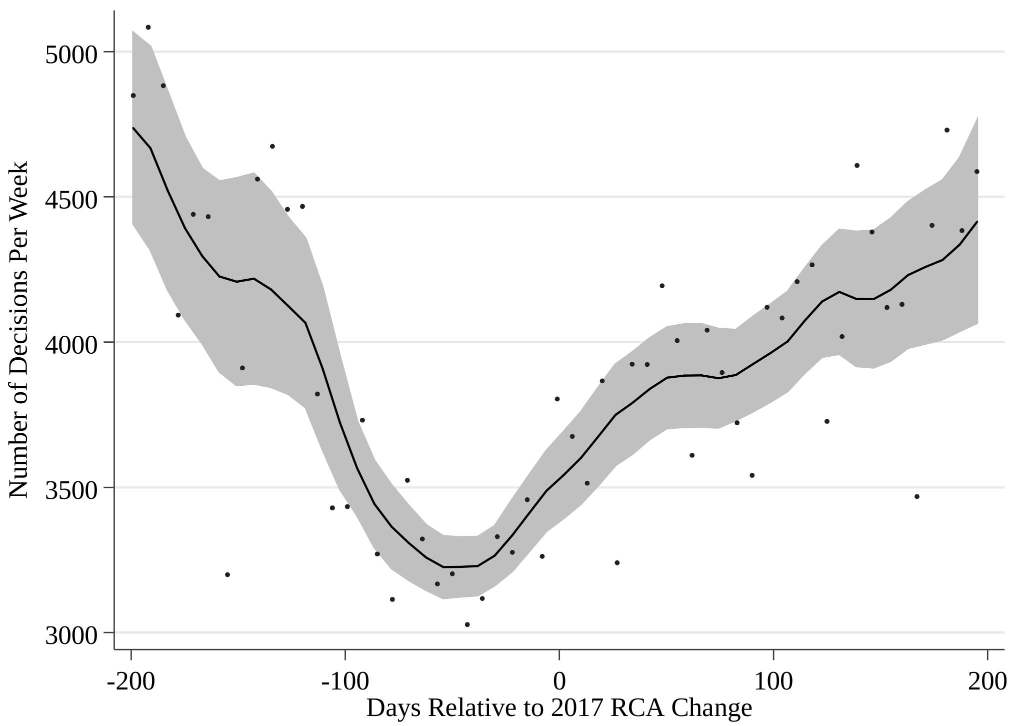

Effects on Immigration Judge Bond Decisions

Finally, we offer a brief qualitative account of how these two changes to the RCA tool affected immigration bond determinations. Recall that, if a noncitizen is not subject to mandatory detention but ICE nonetheless refuses to release him or her, that noncitizen can seek a bond redetermination hearing from an immigration judge.

We draw on data on bond hearings conducted by the Executive Office for Immigration Review to investigate whether these changes in ICE release practices systematically shunted more responsibility onto Immigration Judges ("IJs"). Figure 11 suggests that, when ICE increased its detention rate in 2015, the number of bond hearings jumped, increasing by over 25% in a few months. But IJs did not fully compensate for ICE's change in policy: the number of *grants* of bond did not increase as much as the number of total bond decisions. In other words, ICE's policy change led a large number of people to never obtain release, even though they were able to request it before an immigration judge.

Figure 12 does not show that the 2017 software change had a large effect. That may be partly because the 2017 policy change had a much smaller effect than its 2015 analogue: recall that the 2015 software change was preceded by a large ramp-up in detention, and that the total increase in 2015 was three or four times as large in 2015 as in 2017. Also, the 2017 change occurred just after a much larger policy change: President Trump's dramatic increase in the number of detentions of noncitizens not subject to mandatory detention (see also Ryo and Peacock 2018; Kim and Semet 2019). The software change was small in comparison.

Figure 12: Immigration Judge Bond Hearings and Bond Grants Over Time

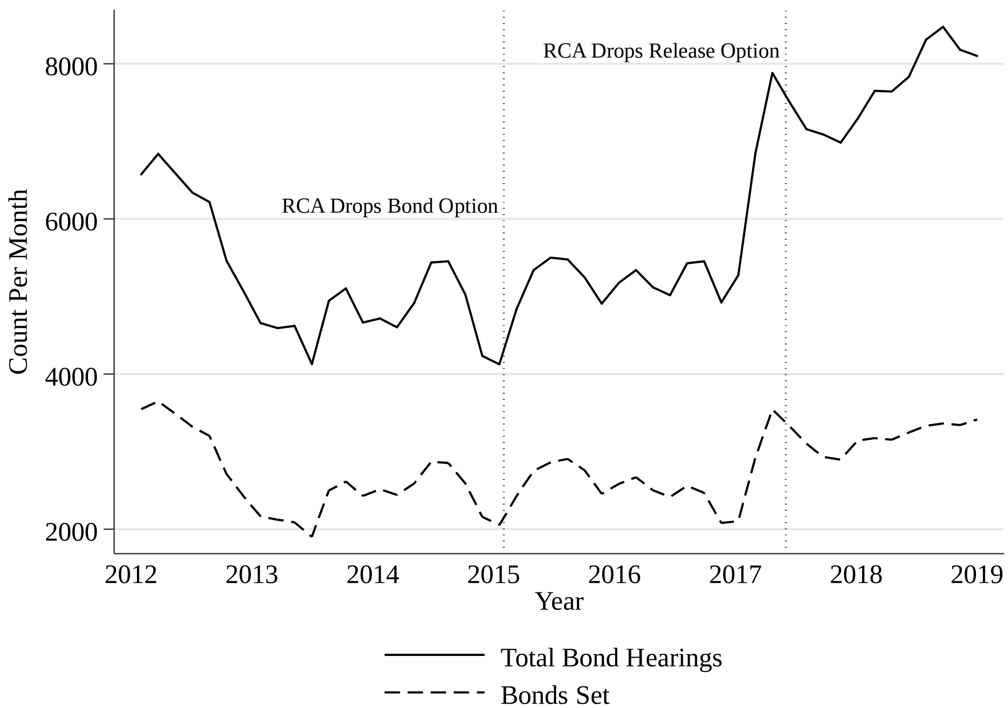

Discussion and Implications

We have shown that ICE manipulated its risk assessment tool to deprive the tool of its ability to make meaningful detention recommendations. That manipulation took place in two steps, in 2015 and 2017. We find, in both 2015 and 2017, that changes to the RCA tool *followed* ICE policy (by detaining more people at a time when ICE had already begun doing so) and, only in 2017, that it *led* that policy (by causing ICE officers to decrease their bond and release rates by about 2-3 percentage points). In the 2015 case, the RCA tool changed to reflect a recent decrease in bond grants; in the 2017 case, the RCA tool changed to reflect a more longstanding ICE policy of detaining more than 90% of people for whom it made risk determinations. Both changes followed steep declines in the number of ICE initial custody determinations, leading us to hypothesize that bed space may have driven ICE's attempts to increase how often it decided to detain those it arrest.

ICE's manipulation of its risk assessment tool has implications for the literatures on algorithmic decisionmaking, immigration detention, and the exercise of executive discretion. These findings highlight a risk of algorithmic decisionmaking tools: that policymakers might modify the tools not to predict risk but instead to implement predetermined policies: we find that ICE changed the RCA tool to match its practicing of detaining more people rather than vice versa. Moreover, we show that the 2017 change to the RCA tool was itself a small—though non-negligible—driver of the growth of immigration detention in the first year of the Trump administration. Finally, these findings highlight ICE's ability nearly to eliminate the discretion of its officers in determining whether to detain noncitizens whom they have arrested; we find no evidence of the type of "hydraulic displacement" of discretion that scholars of criminal proceedings have noted (Miethe 1987).

References

Corbett-Davies, Sam and Sharad Goel. "The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning." Working Paper, 2018.

Corbett-Davies, Sam, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. "Algorithmic Decisionmaking and the Cost of Fairness." KDD Research Paper, 2017.

Eagly, Ingrid V. and Steven Shafer. 2015. "A National Study of Access to Counsel in Immigration Court." *University of Pennsylvania Law Rev*. 32 (2015): 1.

Engen, Rodney L. and Sara Steen. 2000. The Power to Punish: Discretion and Sentencing Reform in the War on Drugs. *American Journal of Sociology* 105 (March 2000): 1357.

Gilman, Denise. 2016. "To Loose the Bonds: the Deceptive Promise of Freedom from Pretrial Immigration Detention." *Indiana Law Journal* 92 (Winter 2016): 157-225.

ICE, Secretary Napolitano and ICE Assistant Secretary Morton Announce New Immigration Detention Reform Initiatives, October 6, 2009, at www.dhs.gov/news/2009/10/06/new-immigration-detention-reforminitiatives-announced.

Kim, Catherine Y. and Amy Semet. 2020."Presidential Ideology and Immigrant Detention." *Duke Law Journal* (forthcoming 2020).

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2017. "Human Decisions and Machine Predictions." National Bureau of Economic Research Working Paper No. 23180 (February 2017).

Koulish, Robert. 2016. "Using Risk to Assess the Legal Violence of Mandatory Detention." MDPI Laws 5 (2016): 30.

Koulish, Robert and Ernesto Calvo. 2019. "The Human Factor: Algorithms, Dissenters, and Detention in Immigration Enforcement." Working Paper, March 6, 2019.

Mayson, Sandra G. 2019. "Bias In, Bias Out." *Yale Law Journal* 128 (2019): 2218.

Miethe, Terance D. 1987. "Charging and Plea Bargaining Practices Under Determinate Sentencing: An Investigation of the Hydraulic Displacement of Discretion." *Journal of Criminal Law and Criminology* 78: 155.

Noferi, Mark and Robert Koulish. 2014. "The Immigration Detention Risk Assessment," *Georgetown Immigration Law J*. 29 (2014): 45.

Ryo, Emily. 2016. "Detained: A Study of Immigration Bond Hearings." *Law & Society Rev*. 50 (2016): 117.

Ryo, Emily & Ian Peacock. 2018. “A National Study of Immigration Detention in the United States.” *Southern California Law Review* 92 (2016): 1.

Ryo, Emily & Ian Peacock. 2019. “Beyond the Walls: The Importance of Community Contexts in Immigration Detention.” *American Behavioral Scientist* 63 (2019): 1250.

Ryo, Emily. 2019. “Predicting Danger in Immigration Courts.” *Law & Social Inquiry* 44 (2019): 227.

Ryo, Emily. 2019a. “Understanding Immigration Detention: Causes, Conditions, Consequences.” *Annual Review of Law and Social Science* 15 (2019): 97.

Schriro, Dora. 2009. Immigration Detention: Overview and Recommendations. Immigration and Customs Enforcement, October 6, 2009. https://www.ice.gov/doclib/about/offices/odpp/pdf/ice-detention-rpt.pdf.

Schriro, Dora. 2019. Expert Report, Hernandez v. Barr. July 29, 2019 (on file with authors).

Ulmer, Jeffrey T., Megan C. Kurlychek, and John H. Kramer. 2007. “Prosecutorial Discretion and the Imposition of Mandatory Minimum Sentences.” *Journal of Research in Crime and Delinquency* 44 (November 2007): 427.

Wilson, Jeffrey. 2019. Deposition in Hernandez v. Barr. March 21, 2019 (on file with authors).

Yang, Crystal S. 2014. “Have Interjudge Sentencing Disparities Increased in an Advisory Guidelines Regime? Evidence from *Booker*.” *New York University Law Review* 89 (2014): 1268.

Appendix: Figures by Field Office (2015 Change on Left, 2017 Change on Right)

Figure 1: Atlanta

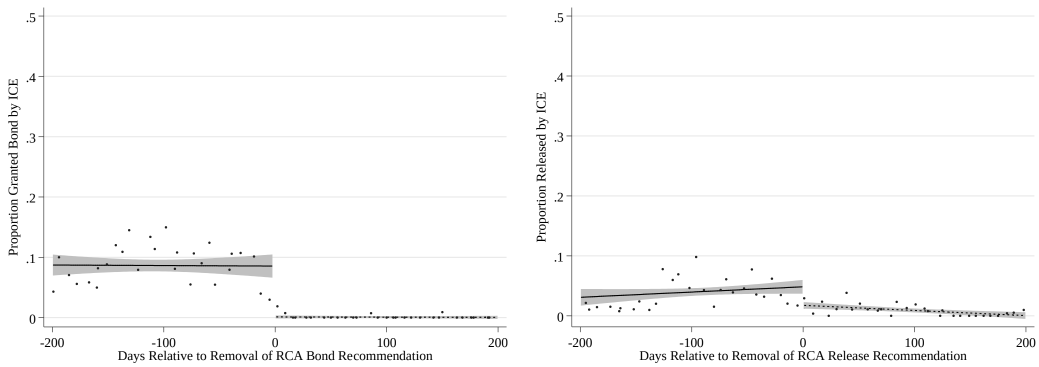

Figure 2: Baltimore

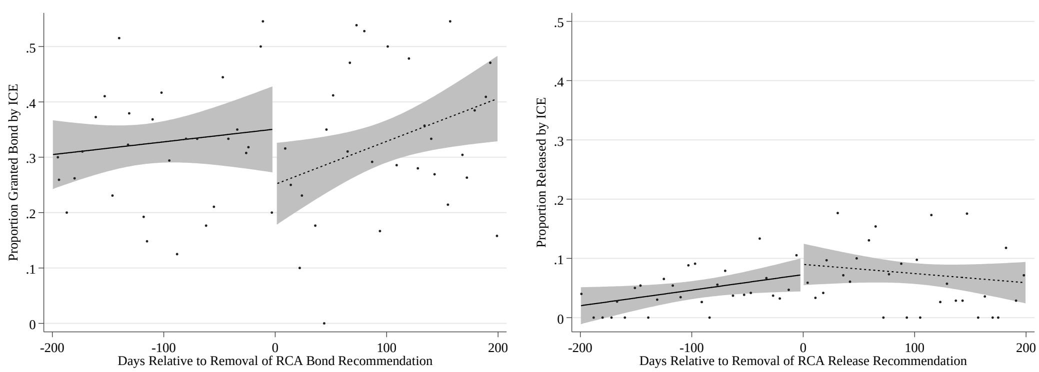

Figure 3: Boston

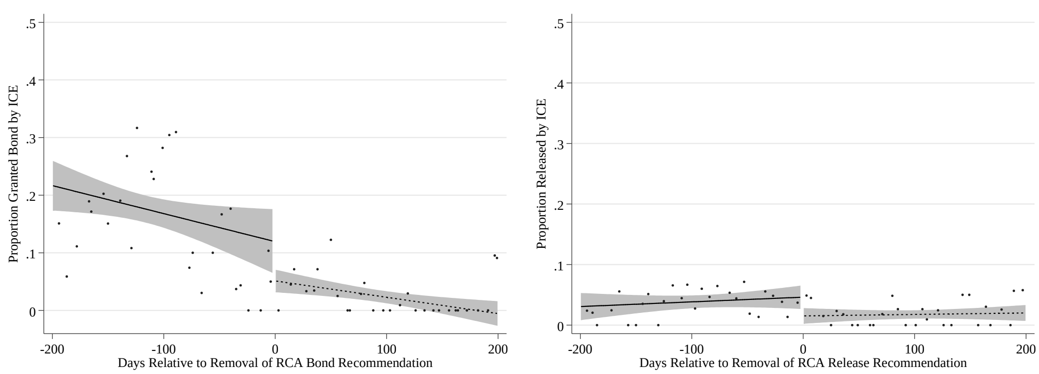

Figure 4: Buffalo

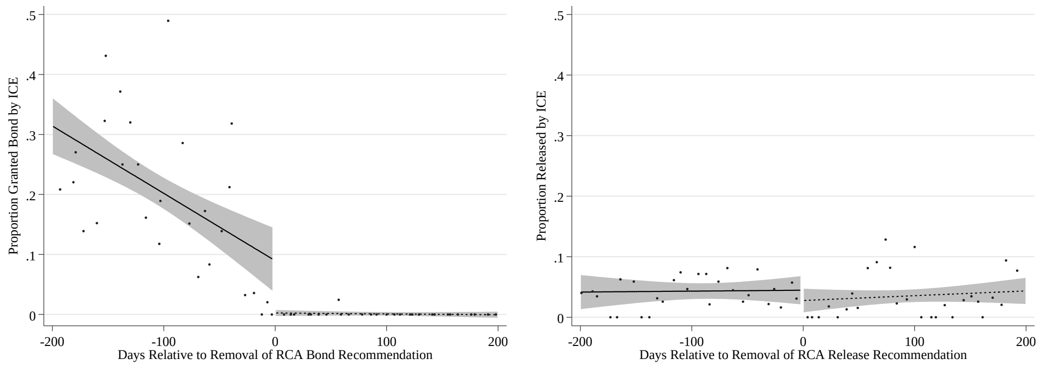

Figure 5: Chicago

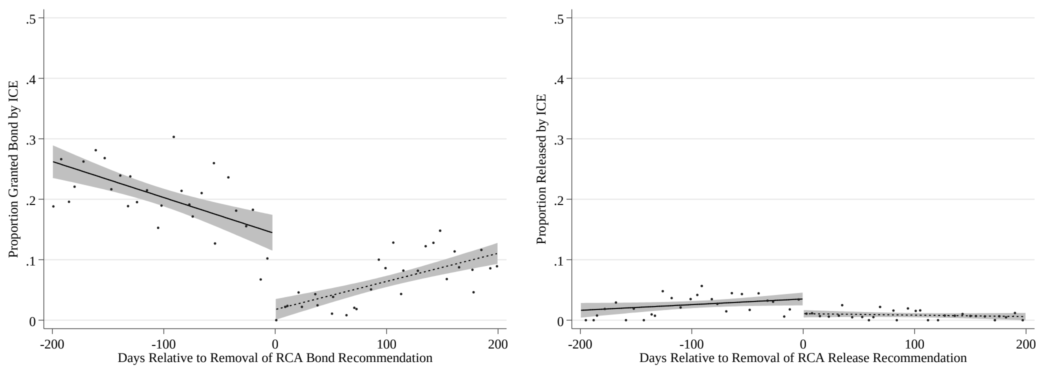

Figure 6: Dallas

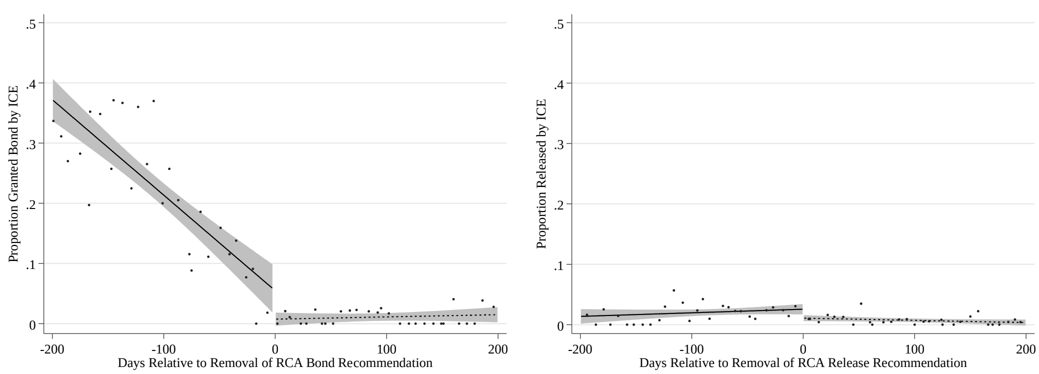

Figure 7: Denver

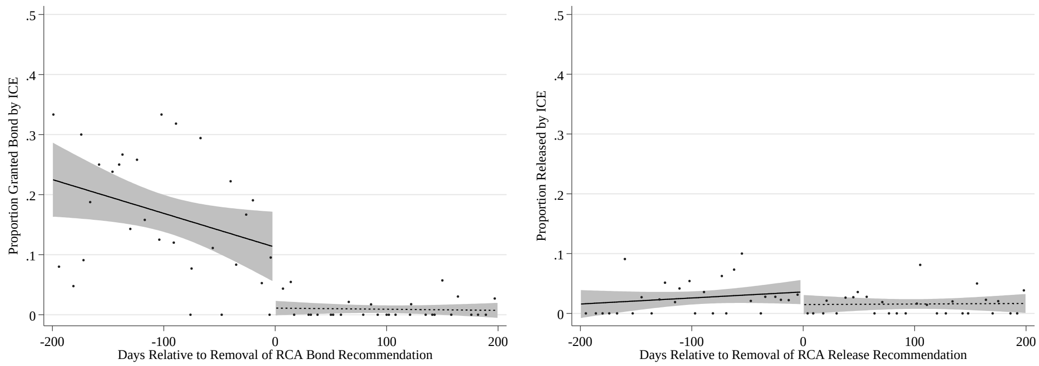

Figure 8: Detroit

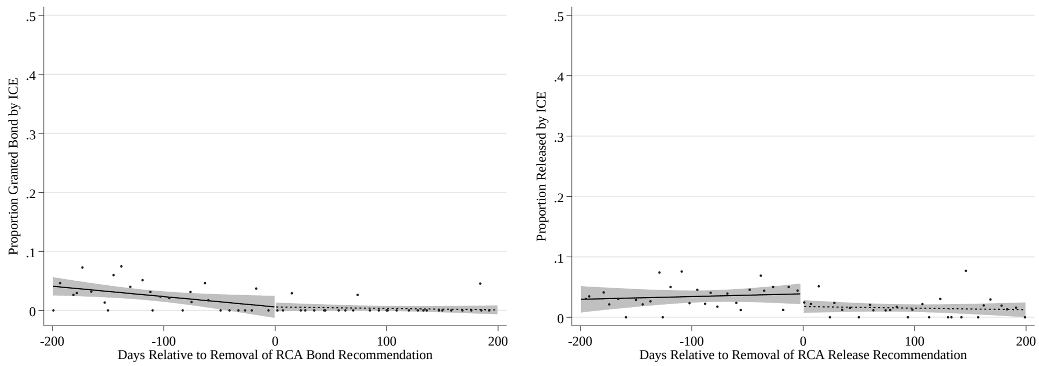

Figure 9: El Paso

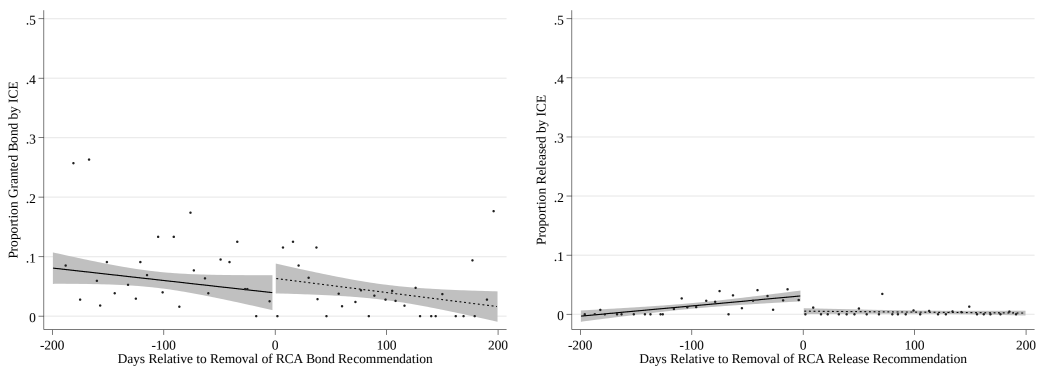

Figure 10: Houston

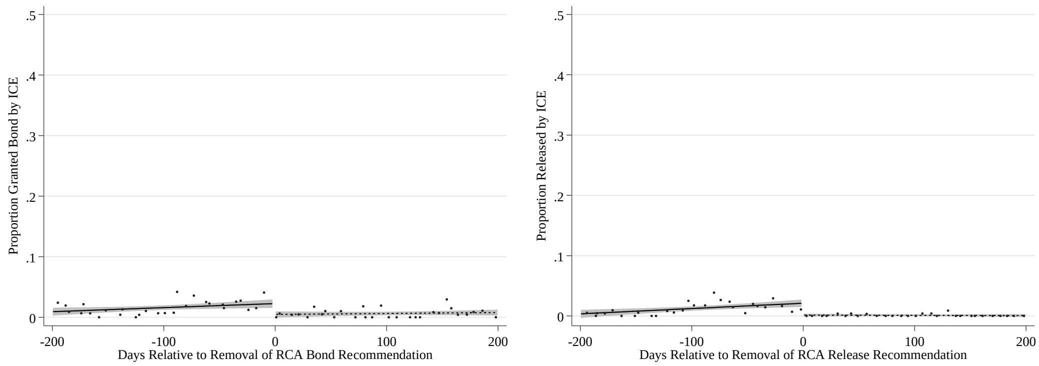

Figure 11: Los Angeles

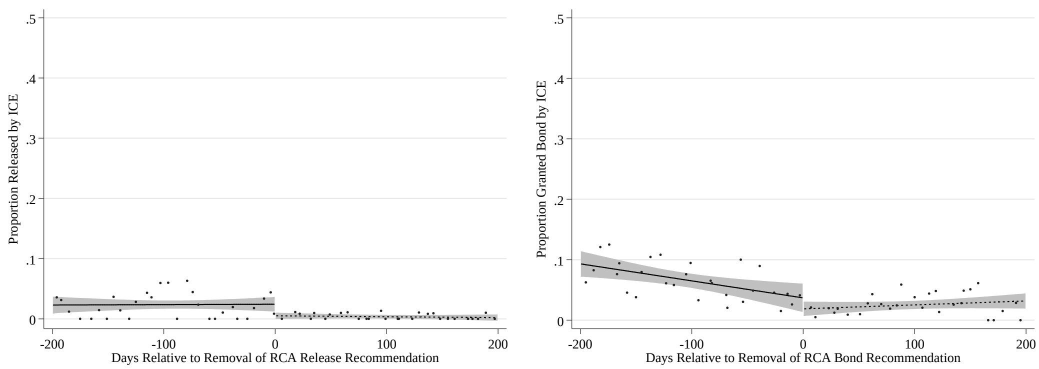

Figure 12: Miami

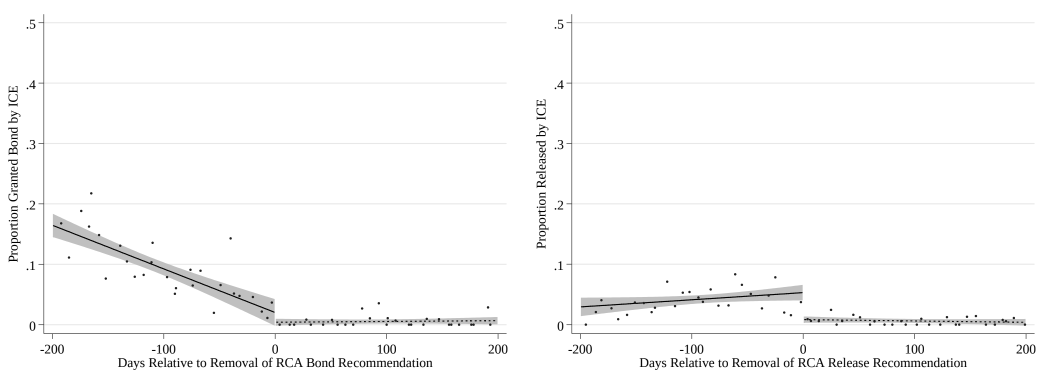

Figure 13: New Jersey

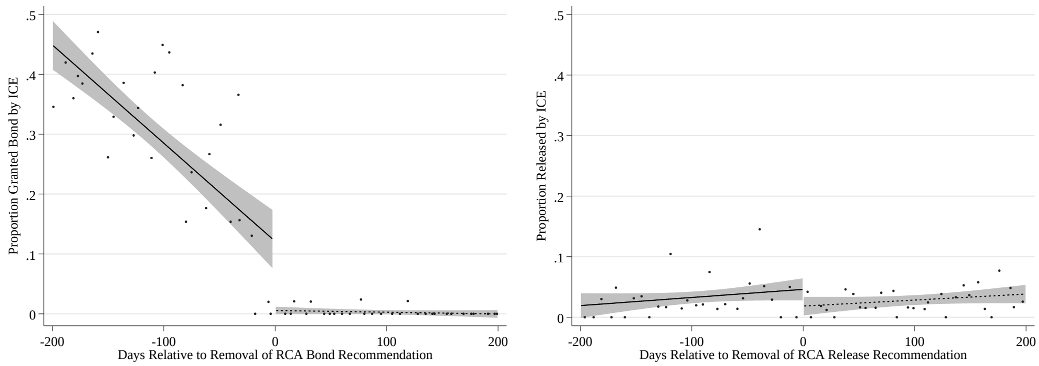

Figure 14: New Orleans

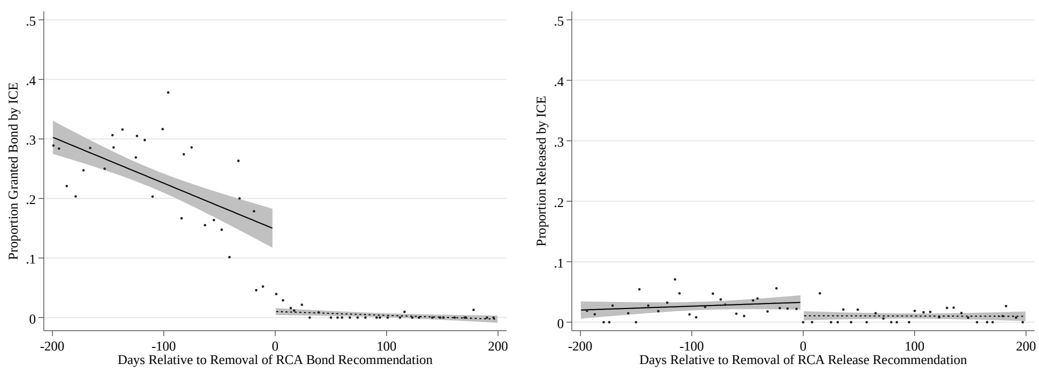

Figure 15: New York

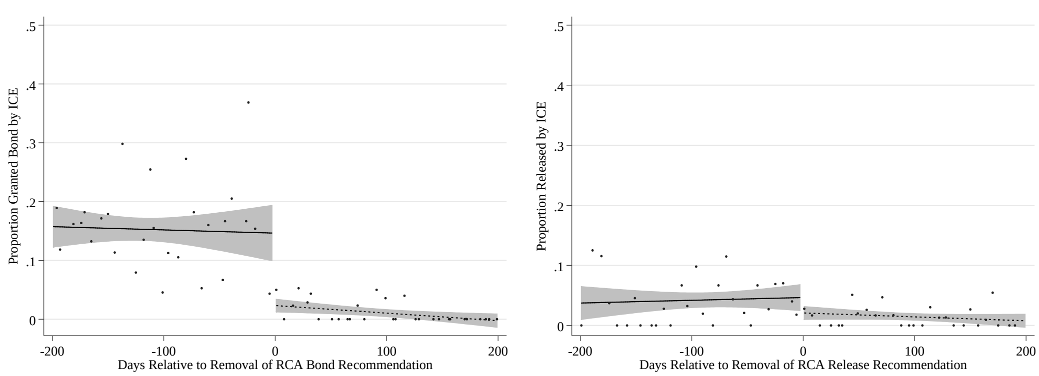

Figure 16: Philadelphia

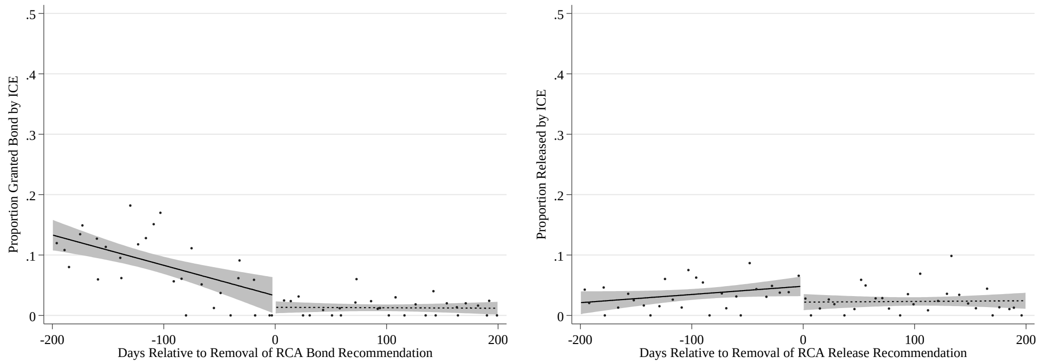

Figure 17: Phoenix

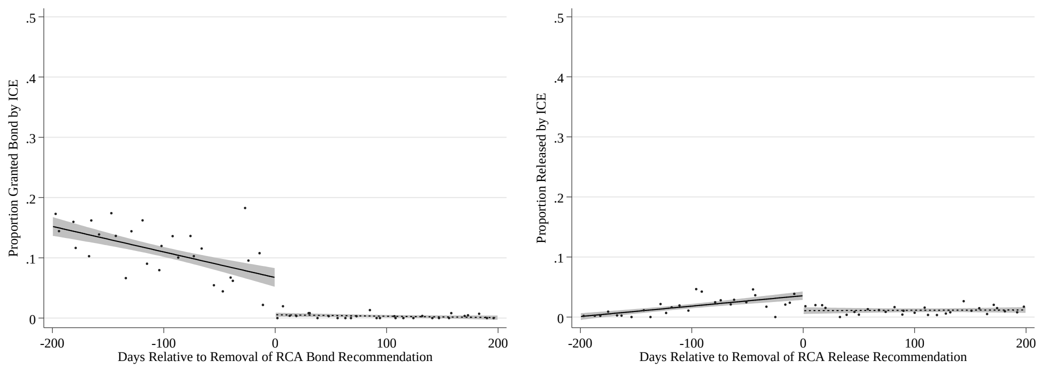

Figure 18: Seattle

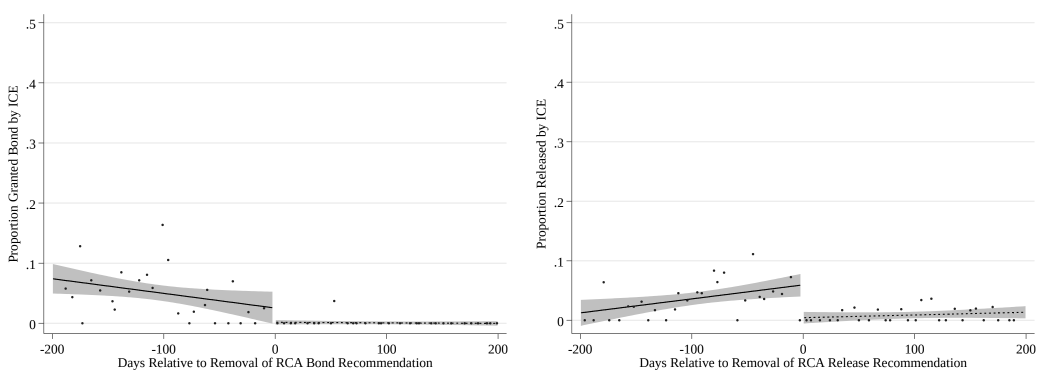

Figure 19: San Francisco

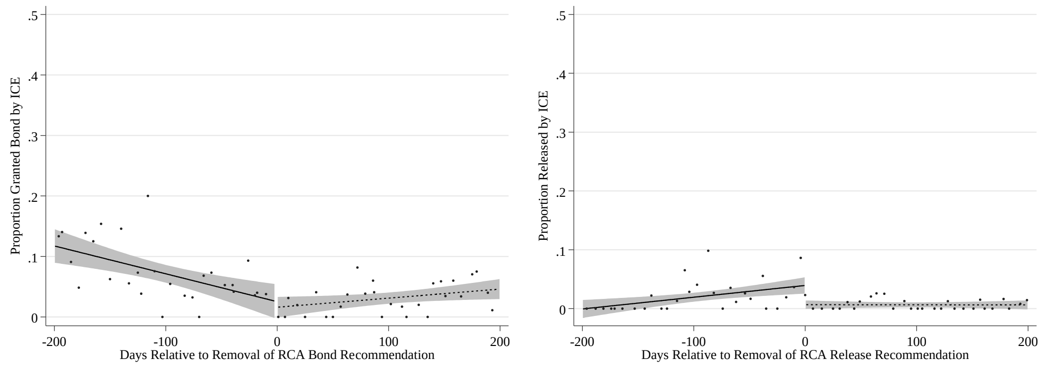

Figure 20: Salt Lake City

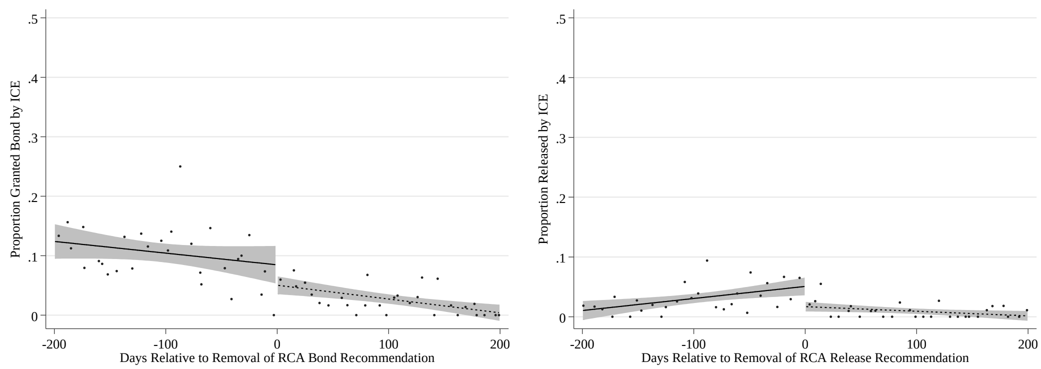

Figure 21: San Antonio

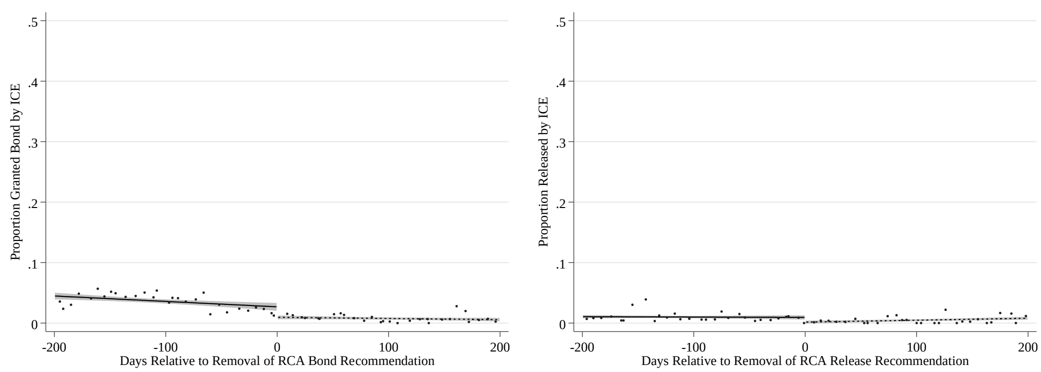

Figure 22: San Diego

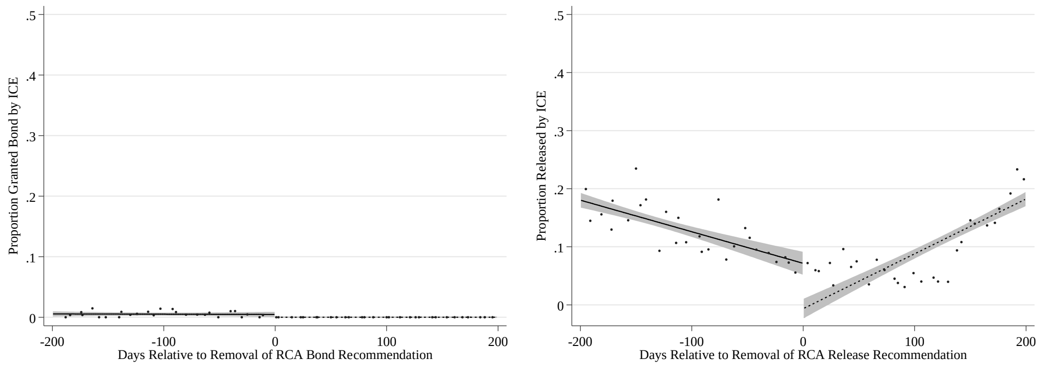

Figure 23: Saint Paul

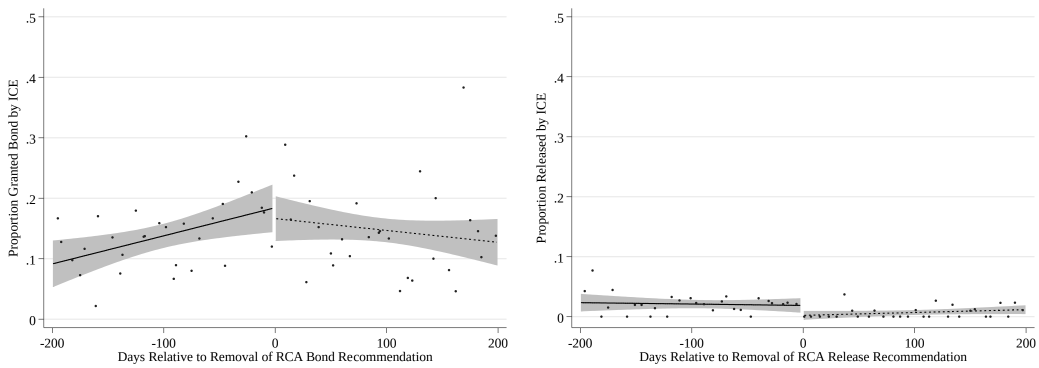

Figure 24: Washington