Restraining Orders and the Deterrence of Domestic Violence*

Eric Helland†
Claremont McKenna College and RAND

Minjae Yun‡
Claremont McKenna College

July 20, 2026

Abstract

This study estimates the causal effect of civil restraining orders on subsequent crime using administrative records on restraining-order petitions filed in Los Angeles between 2020 and 2023 linked to LAPD crime reports. Exploiting quasi-random variation in judicial propensity to grant petitions, we implement a judge-leniency instrumental variables design. We find little evidence that granting a restraining order reduces crime when all petition types are pooled. However, substantial treatment heterogeneity emerges across order types. Granting a domestic violence restraining order involving minor children reduces the probability of a subsequent police-reported domestic violence incident by 33–41 percentage points, while we find no comparable deterrent effect for civil harassment petitions or domestic violence petitions without minor children. These findings suggest that restraining orders should not be viewed as a homogeneous intervention; their effectiveness depends on the broader bundle of family-law remedies accompanying the order.

*The authors wish to thank Paul Heaton, Aurélie Ouss, Jennifer Doleac, Aviv Caspi, Yotem Shem-Tov, Ben Barsky, and participants at ALEA (NYU) and CELS (Georgetown) for their helpful comments. They also appreciate the assistance of several judicial officers and court personnel at LASC, both for their valuable insights and critical guidance in understanding family court procedures generally and, specifically, the process in the stand-alone restraining order courts. We also note that our research was produced independently and reflects the personal views of its authors and not the views of the Los Angeles Superior Court or any other entity.

†Robert Day School of Economics and Finance, Claremont McKenna College, 500 East Ninth Street, Claremont, California, 91711 (eric.helland@cmc.edu).

‡Robert Day School of Economics and Finance, Claremont McKenna College, 500 East Ninth Street, Claremont, California, 91711 (minjae.yun@cmc.edu).

1 Introduction

The CDC estimates that the lifetime prevalence of intimate partner violence (IPV), including physical violence, sexual violence, and stalking by an intimate partner, exceeds one in three women (35.6%) and one in four men (28.5%) (Black et al., 2011). Moreover, victims frequently experience repeated abuse (Aizer and Dal Bo, 2009). In Los Angeles County, the location of our study, a 2010 survey by Blue Shield of California similarly found that one in five respondents had experienced domestic violence at least once during their lifetime (see Los Angeles County Department of Public Health, Injury and Violence Prevention Program, 2013). The consequences of IPV are often severe. In Los Angeles County, an intimate partner kills a woman approximately every five days (Tjaden, 2000). Beginning in the 1970s, states responded to the inadequacies of existing legal remedies with several major policy reforms, including mandatory or presumptive arrest policies, no-drop prosecution policies, and the adoption of civil restraining orders that provided victims with a civil, rather than criminal, mechanism for preventing future contact with an alleged abuser.[1]

Civil restraining orders are intended to provide immediate protection by legally prohibiting future contact between the respondent and the petitioner. In California, as in most states, violating a restraining order is itself a criminal offense, allowing law enforcement to arrest and prosecute respondents who fail to comply with the order, although arrests based solely on restraining-order violations are relatively uncommon.[2] Given the prevalence of IPV, a substantial literature examines policies designed to reduce subsequent violence. Most of this work has focused on criminal justice interventions, particularly mandatory arrest policies and no-drop prosecution.[3]

By contrast, relatively little research has examined the effectiveness of civil restraining orders despite their widespread use and their role as the primary civil remedy available to victims of domestic violence (Ko, 2001). Early studies generally relied on interviews with women in domestic violence shelters or descriptive comparisons of petitioners who retained and dropped their restraining orders.[4] A small number of studies linked court records to subsequent arrests. For example, Klein (1996) examined 663 restraining-order cases in Quincy, Massachusetts, and found no significant differences in subsequent abuse between petitioners who retained and those who dropped their restraining orders. In contrast, Carlson et al. (1999) used court

[1] In 1976, only two states had laws allowing civil restraining orders. By 1994, all 50 states had adopted some form of protective-order statute, often accompanied by specialized courts dedicated to hearing these petitions (Ko, 2001).

[2] Civil protection orders are, in some ways, an alternative to criminal charges because they are immediately available to eligible victims and often provide broader protection than criminal protective orders, which generally require an arrest before issuance (see Ko, 2001; Judicial Council of California, 2023).

[3] Mandatory arrest policies have been studied extensively, beginning with a series of randomized experiments (see Sherman and Berk, 2017; Dunford et al., 1990; Hirschel and Hutchison iII, 1992; Sherman et al., 1992; Pate and Hamilton, 1992; Berk et al., 1992). More recent work has used difference-in-differences, event studies, and instrumental variables to study these policies. For example, Iyengar (2009) finds evidence that mandatory arrest policies increase violence, while Chin and Cunningham (2019) finds no evidence that mandatory arrest laws increase spousal homicides and some evidence that discretionary arrest laws reduce them. Using a judge-leniency design similar to ours, Amaral et al. (2023) finds that IPV arrests reduce future domestic violence. No-drop prosecution policies are examined by Aizer and Dal Bo (2009).

[4] Studies examining the physiological benefits of restraining orders are more common but face many of the same methodological limitations (see Ko, 2001).

and police records and found substantial reductions in domestic violence arrests following the issuance of permanent restraining orders, with larger effects for permanent than temporary restraining orders.

Although these studies provide important descriptive evidence, they face two important limitations. First, most rely on observational comparisons without a credible control group, making it difficult to distinguish the effect of the restraining order from pre-existing differences across petitioners. Second, even studies using administrative data cannot readily separate the causal effect of granting a restraining order from the fact that judges are more likely to grant petitions when they believe respondents pose a greater threat of future violence. Petitioners whose requests are granted, therefore, differ systematically from those whose requests are denied. More fundamentally, the existing literature implicitly treats restraining orders as a homogeneous intervention, despite substantial institutional differences across order types. Our objective is to overcome these identification problems by exploiting quasi-random variation in judicial leniency to estimate the causal effect of granting a restraining order on subsequent crime and to examine whether those effects differ across the distinct legal interventions that California groups under the umbrella of "civil restraining orders."

Our primary data source consists of 20,384 restraining-order petitions filed in the Los Angeles Superior Court between 2020 and 2023. The case information was scraped from the court's online case-access system and supplemented with administrative records provided by the Los Angeles Superior Court (LASC), including the block-level address protected by the requested order.[5] Because our outcomes are measured at the block level, we restrict the analysis to petitions that identify a protected residential address. Approximately 20% of petitions, primarily civil harassment cases requesting only "stay-away" provisions, do not contain a residential address and are therefore excluded from the analysis.[6]

The treatment in our study is the judicial decision at the mandatory 21-day hearing. In the majority of cases, petitioners have already received a temporary restraining order pending that hearing. A judge's decision therefore determines whether the temporary order expires or is replaced by a longer-term restraining order, generally lasting between three and five years. Of the 20,384 petitions in our sample, judges granted 4,803 and denied the remainder, providing the treatment and comparison groups used throughout the analysis.

We measure subsequent victimization by matching the protected address in each petition to the Los Angeles Police Department's block-level crime database. Because both the court records and the crime data identify locations only at the block level, the unit of observation is the residential block listed in the petition. Our primary outcome is an indicator for whether a reported residential domestic violence incident occurred on that block during the year following the 21-day hearing. We also consider broader measures of assault as alternative outcomes.

[5]To validate the court-provided addresses, we randomly selected 4,000 petitions and manually extracted the protected addresses from the underlying court filings. Among petitions containing a residential address, only approximately 1.5% differed at the block level from the administrative data.

[6]These cases typically involve stalking or harassment occurring away from the petitioner's residence, making a residential crime measure infeasible.

Our empirical strategy exploits the as-if random assignment of judges to the mandatory 21-day restraining-order hearing. Because judges differ systematically in their propensity to grant petitions, judicial assignment generates plausibly exogenous variation in whether a petitioner receives a long-term restraining order. We use this variation in a judge-leniency instrumental variables design to identify a Local Average Treatment Effect (LATE): the causal effect of granting a restraining order for the subset of petitions whose disposition is influenced by judicial assignment, independent of the selection bias present in ordinary least squares estimates (see Agan et al., 2023; Chyn et al., 2025).

Using this design, we find little evidence that granting a restraining order reduces subsequent crime when all petitions are pooled. Disaggregating the analysis, however, reveals substantial treatment heterogeneity. Granting a restraining order reduces the probability of a subsequent police-reported residential domestic violence incident by approximately 26-40 percentage points for domestic violence restraining order (DVRO) petitions involving minor children. In contrast, we find no comparable deterrent effect for civil harassment petitions or domestic violence petitions without minor children, despite strong first stages and comparable statistical power across petition types.

This pattern is surprising because the criminal sanctions for violating a restraining order are largely the same across the different categories of civil restraining orders.[7] The institutional consequences of granting an order, however, differ substantially across petition types. In particular, domestic violence restraining orders involving minor children frequently bundle additional family-law remedies, including custody, visitation restrictions, exclusive possession of the residence, and temporary support orders, that substantially reduce the parties' continued interaction. These accompanying legal remedies suggest that restraining orders should be viewed not as a single intervention but as distinct legal bundles whose effectiveness depends on the institutional context in which they operate.

This paper makes three contributions to the literature on civil restraining orders. First, we provide causal estimates of the effect of granting a restraining order on subsequent crime using a judge-leniency research design. Second, we show that restraining orders are highly heterogeneous interventions, with meaningful deterrent effects concentrated among domestic violence petitions involving minor children. Third, we provide evidence that this heterogeneity reflects the broader bundle of legal rights and obligations accompanying these orders rather than the criminal sanctions associated with violating the restraining order alone.

The remainder of the paper proceeds as follows. Section 2 describes the legal framework governing civil restraining orders and its implications for our research design. Section 3 discusses the data. Section 4 presents the empirical strategy. Section 5 reports the main results and robustness analyses. Section 6 concludes.

[7] Both civil harassment orders and domestic violence restraining orders are misdemeanors punishable by up to one year in county jail and a $1,000 fine under California Penal Code §273.6, with similar aggravating factors that may elevate violations to felonies.

2 Background

In this section, we offer some context for civil restraining orders in California before turning to a discussion of the process for obtaining a civil restraining order in Los Angeles, in which petitioners can obtain a civil restraining either in a special court focusing on restraining orders (stand-alone restraining order courts), which are the focus of our study, or in family court as part of a broader case such as divorce.

2.1 Civil Restraining Orders in California

Both civil and criminal courts can issue restraining orders.[8] The procedures and criteria for obtaining these orders differ across jurisdictions based on specific legal statutes. In California, protective orders are typically issued during criminal proceedings at the judge's discretion, typically when the perpetrator has been arrested for domestic violence. These orders, which may be a condition for bail or pretrial release, instruct the offender not to threaten, harm, or assault the victim while the case is ongoing, such as during arraignment or sentencing. Generally, protective orders are temporary and remain in effect only until the criminal case concludes, at which point the court's jurisdiction over the defendant ends. However, judges have the authority to modify or terminate the orders earlier. Violating a protective order can result in additional penalties, including fines or jail time. In California, violating such an order can result in charges ranging from a misdemeanor to a felony. Generally, criminal protective orders seem to be viewed by the courts as part of the prosecution process rather than a primary means of safeguarding victims in contrast to civil restraining orders (see Ko (2001) and the Judicial Council of California (2023)).[9] Violations of civil restraining orders can also carry a misdemeanor or felony charge.[10]

In California, requests for restraining orders are typically handled in family court. There are several different restraining orders, the most common being the Domestic Violence Restraining Order (DVRO), followed by, in order of frequency, civil harassment, elder abuse, gun violence, workplace violence, and school violence (see Table 1). Domestic violence restraining orders are filed when abuse is perpetrated by a person with whom the petitioner has or once had a relationship, such as a spouse, partner, or close relative.[11] Civil harassment restraining orders are filed against people with whom the petitioner has no close relationship, such as neighbors, coworkers, or non-immediate relatives. Elder or dependent adult abuse restraining orders are filed to protect elders or dependent adults who are abused or neglected, and they can be filed by someone

[8] In criminal cases, they are often termed "protective orders" or "stay-away orders," while in civil cases, they are typically called "civil restraining orders" or "civil protective orders." However, the terminology and definitions of these orders vary by state, with some states using the terms interchangeably.

[9] There is an overlap between criminal and civil protective or restraining orders since a petitioner for a civil restraining order could also be protected by a criminal order. This overlap is noted on the civil order; however, in our sample, we did not identify any restraining order requests that also had a criminal protective order.

[10] Anecdotally, standalone enforcement of civil restraining orders is rare, and typically, penalties for violations of civil restraining orders arise out of charges for other criminal offenses against the petitioner by the respondent (see Harper et al. (2021).)

[11] Abuse does not have to be the actual infliction of physical injury or assault for a restraining order to be granted.

else on behalf of the elder or dependent adult. The three remaining types are specific to certain occupations. Police officers or sheriffs typically file gun violence restraining orders. Employers file workplace violence restraining orders on behalf of an employee, while school representatives file school violence restraining orders on behalf of a student.[12]

2.2 Obtaining a Restraining Order

The process of obtaining a civil restraining order in Los Angeles County is a relatively straightforward two-step process. The first step is determining the appropriate type of restraining order based on the nature of the relationship between the petitioner and the respondent. Petitioners must complete three specific forms requesting the order and, if applicable, a temporary order to cover the period until the restraining order hearing occurs. These forms detail the incidents of abuse or harassment, and the type of protection sought. Once completed, the forms are filed with the court in the geographic area where the incident occurred or where the petitioner resides. Although there are filing fees for some orders, fee waivers are available for cases involving violence, threats, or stalking, and no filing fee is required for Domestic Violence Restraining Orders. In the request, the petitioner describes past incidents of abuse and identifies any children, family members, or cohabitants they want the restraining order to protect. They must also state whether the respondent possesses any firearms. Additional forms may be required if the petitioner seeks child custody, eviction of the respondent from a shared residence, or spousal support as part of the order request.

Los Angeles County has specialized restraining order courts at certain courthouses.[13] The largest and the focus of our study is at Stanley Mosk Courthouse in downtown Los Angeles, which handles restraining orders originating from residents of the city of Los Angeles. Although each courthouse typically serves its local jurisdiction, the restraining order courts do not have exclusive jurisdiction over all restraining order requests within their geographic area. However, orders outside of the geographic jurisdiction are typically forwarded to the courthouse with the appropriate jurisdiction. In our sample, all cases originated at Stanley Mosk, covered residents of the city of Los Angeles, and were heard at Stanley Mosk. Each courthouse with a restraining order court has a legal aid clinic, typically staffed by attorneys. Although these clinics do not represent the parties, they can assist in selecting the appropriate order category to request, completing forms, and navigating the process.

Requests for temporary restraining orders (TROs) can be made without notifying the other party (ex parte).

[12] One final type of order not included in our sample is an Emergency Protective Order (EPO), which can be issued if there is an immediate threat. A law enforcement officer must request these orders, which typically occur when an officer is responding to a domestic violence call or another emergency call. An EPO lasts 5-7 days and typically prohibits the respondent from possessing firearms or ammunition. Note that the firearm restrictions can also apply in DVROs if the petitioner indicates to the court that the respondent has a firearm.

[13] The specialized LASC restraining order courts are Stanley Mosk Courthouse (the city of Los Angeles), Norwalk Courthouse, Pasadena Courthouse, Pomona Courthouse, San Fernando Courthouse, Santa Monica Courthouse, Torrance Courthouse, and Van Nuys East Courthouse. Given the broad geographic distribution, most orders are filed in the courthouse nearest the petitioner's residence.

The judge typically reviews the TRO request the same day unless it is submitted too late, in which case the decision will be issued the next day. A TRO can include orders to prevent abuse, move out, stay away, control property, provide spousal and child support, and determine child custody and debts. The threshold for a TRO and the restraining order after hearing (ROAH) is "reasonable proof of past acts of abuse," and TROs are issued at the discretion of the judicial officer when "good cause" is shown.

Judges grant about 80% of the TRO requests. The difficulty in using the TRO period in the analysis discussed below is that some TROs do not go into effect due to a lack of service.[14] Other TROs do not go into effect because the petitioner does not prosecute the request after filing the initial forms or because the petitioner withdraws the request sometime after the TRO is granted.[15] If the TRO is granted, it remains valid until the 21-day hearing. In domestic violence cases, the TRO automatically includes a prohibition of firearm possession by the respondent and restrictions on searching for protected persons. If the court grants a continuance, the TRO remains in effect until the hearing concludes unless otherwise directed by the court. Regardless of the decision on the TRO, the court must schedule a hearing to consider the full request within 21 days or 25 days if good cause for a delay is shown. By law, restraining order hearings take precedence over all other calendar events. Due to persistent staffing issues in the specialized restraining order courts, the 21-day requirement often results in restraining order hearings being referred to general judges at Stanley Mosk when commissioners are unavailable for the required 21-day post-filing hearing (see below).

At the restraining order hearing, the judge decides whether to grant a long-term restraining order, known as a Restraining Order After Hearing (ROAH). The court considers the affidavit, testimony, and any additional information. If they attend, the respondents can contest the claims of abuse. As with the TRO, the standard is the preponderance of the evidence.[16]

Broadly, there are four reasons why TROs or ROAHs are denied. The first is that the order request does not belong in the restraining order court and should have been filed in another area of family court (e.g., the petition does not allege abuse but is seeking a custody arrangement). The second is that the petition does not attend the 21-day hearing (failure to appear (FTA)), and no other evidence or testimony on the alleged abuse was provided to the court; the third, typically conditional on the petitioner attending the hearing, is because the evidence is not strong enough to meet the threshold; and finally, order requests are more likely to be denied if the respondent contests them.[17]

[14]The respondent must be served at least five calendar days before the 21-day hearing. While law enforcement can be used to notify the respondent of the TRO, any adult over 18 years of age can serve the documents. However, it is the petitioner's responsibility to arrange for service and notify the court of proof of service. The TRO is not effective until served. By contrast, ROAHs are not approved without proof of service.

[15]For our estimation, this means that it is very difficult to determine when the TRO was in effect, in contrast to ROAHs, in which the date of the 21-day hearing effectively continues or ends the restraining order.

[16]Note that petitioners do not have to attend the 21-day hearing to have their request granted, although, in practice, the petitioner's testimony is typically critical to the decision to grant a restraining order. A single testimony, even from the petitioner, can be sufficient to issue a DVRO. The time since the most recent abuse does not solely determine the outcome; the court evaluates the totality of circumstances. A likelihood of future abuse is not required to issue a restraining order. It is worth noting that this standard is weaker than that in criminal cases, which require proof beyond a reasonable doubt.

[17]One concern is that petitioner's failure to appear may be driven by court logistics that could undermine randomization.

Judges may also rule at the hearing on issues such as child and spousal support, legal fees, batterer intervention programs, etc. Unlike TROs, restraining orders can evict the respondent from housing shared with the petitioner. An ROAH can last up to five years. If no expiration date is specified, the order remains valid for three years.

Finally, and importantly for our leniency design instrumental variables strategy, the hearing ends the judge's involvement in the case unless the parties request a renewal or modification. It is essentially impossible for a judge to impact the likelihood of an arrest beyond issuing the order. LASC does not, for example, have a method beyond the petitioner's sworn affidavit to determine if the firearm was surrendered or if the respondent was truthful that they did not have a firearm to begin with. If the ROAH is granted, the terms of the order are entered into the California Law Enforcement Telecommunications System (CLETS), a law enforcement database that provides information about restraining orders issued by the court.[18]

2.3 Judge Assignment for the 21-Day Hearing

As mentioned above, the Los Angeles Superior Court (LASC) has several standalone restraining order courts. At Stanley Mosk during our sample period, these courts were staffed by 6–8 judicial officers, most often commissioners rather than judges.[19] When the petition is initially filed, the court randomly assigns a commissioner to rule on the ex parte temporary restraining order (TRO) request. These determinations are typically made on the filing rather than in person. When the 21-day hearing arises, the commissioners of the standalone restraining order courts are randomly assigned to preside. If a commissioner is not available on that day, or, rarely, if one of the parties objects to having a commissioner hear the case, the matter is reassigned to a judge sitting in the family law division. In the event that no family law judge is available, a judge from the general civil division is randomly assigned.

This assignment process creates two issues for constructing the JIVE estimator used in our judge/examiner IV design discussed below. The first is that randomization occurs at the calendar-day level. Chyn et al. (2025) emphasize that the JIVE estimator should be residualized at the level of randomization. In our context, if the hearing falls on August 30, any commissioner in the standalone restraining order court, or

Specifically, Graef et al. (2023) find that much of the observed FTA is not the defendants who failed to appear, but other parties that disrupted the case. Nicely, these concerns do not apply in our setting. In restraining order cases, the 21-day statutory hearing deadline forces LASC to prioritize these hearings, as we discuss below, ensuring that scheduling logistics do not compromise the random assignment of judges. Essentially, the hearing will take place regardless of whether the petitioner, respondent, lawyers, etc., appear. A second concern, addressed below, is that the timing of the hearing may impact both the availability of judges assigned to the stand-alone restraining order courts and the petitioner FTA. To control for this, we explicitly residualize our instruments on day-of-week, week-of-year, and year fixed effects, and in an alternative specification on calendar date.

[18] CLETS is a secure statewide database used by law enforcement agencies to access information about restraining orders, arrest warrants, criminal records, and other critical data. When the court issues a restraining order, it is entered into the CLETS system, making it accessible to law enforcement officers across the state for enforcement and verification purposes.

[19] In the Los Angeles Superior Court, both judges and commissioners preside over civil case calendars. Judges are elected or appointed judicial officers, while commissioners are subordinate judicial officers appointed by the court who serve at the pleasure of the judges. Commissioners may hear and determine matters only when the parties stipulate to their authority, but in restraining order cases, such stipulations are typically provided as a matter of course. In practice, commissioners perform functions similar to judges in these settings.

any available family or civil judge serving on that day, could potentially hear the case. To account for this, we residualize the grant probability using day-of-week, month-of-year, and year fixed effects. In a separate specification, we estimate the model with calendar-date fixed effects for every hearing date in the sample.

Unlike Dobbie et al. (2018) and Chyn et al. (2025), who can construct the set of all judges available on the evening of a bail hearing from their sample of observed hearings, the relevant availability set in our setting is more complex. Theoretically, any commissioner in the standalone restraining order court, any family law judge, or any general civil judge serving on that day could be assigned. Fortunately, LASC shared with us the roster of judges available on each day of our sample period. We therefore construct the JIVE instrument using only the set of judicial officers available on a given hearing date, even if some were not ultimately selected to hear a restraining order request on that day.

Our second issue is that some judges, i.e, judges relatively new to family court or general civil judges, hear relatively few cases. We exclude any judge who heard fewer than five cases in the year prior to a hearing. Given our limited sample size and the frequency with which cases are referred out of the standalone courts, we cannot raise this exclusion threshold further. To address the resulting imprecision, we follow Agan et al. (2023) and report Anderson–Rubin 95% confidence intervals (Anderson and Rubin, 1949), which provide valid inference even when some first-stage estimates are weak.

Finally, we construct the JIVE estimator using only hearings the judge conducted prior to the current 21-day hearing, or in an alternative specification, using outcomes from the preceding year only. Some studies use the full sample of data (both past and future decisions, excluding the current case) to construct leave-one-out measures. In our setting, however, family law judges regularly receive training on handling restraining order cases, and such training may change their behavior in ways that introduce non-stationarity. Because we do not observe the timing of training or even if it has occurred, we elect to use retrospective behavior only when constructing our JIVE estimator.

3 Data

3.1 Restraining Order Data

Our primary data source on restraining orders contains information on 20,384 restraining order requests filed in the city of Los Angeles between 2020 and early 2023. These data were scraped from the Los Angeles Superior Court's (LASC) online case access system for family court.[20] We matched the scraped data to block-level address information from LASC. LASC's address information is collected to enable the court to correspond with the petitioner, and therefore, it is possible that the address is not the address from which the respondent would be barred. To verify the frequency with which the correspondence address differs from the address from which the petitioner would be banned if the order is approved, we randomly selected

[20] `https://www.lacourt.org/casesummary/ui/index.aspx?casetype=familylaw`

4,000 cases from the initial dataset and examined the restraining order requests. Among these, 3,532 orders included a specific residential address from which the respondent, the alleged abuser, would be barred if the order was approved.[21]

Restraining order requests that lacked a specific address were excluded from our analysis because the nature of our estimation strategy requires crimes to be observed at specific block locations (see below), so orders that require only that the respondent "stay away" from the petitioner cannot be analyzed with our crime data. We keep a sample of 20,384 requests for restraining orders. Among these, 4,803 requests were granted, and 15,328 were denied. The denied requests serve as the control group in our analysis.

Almost all remaining requests initially had a temporary restraining order (TRO) in place, which, if service was conducted, would continue until the hearing. Thus, as mentioned above, the treatment in our study reflects whether an individual moved from having a temporary restraining order to continuing that order for 3-5 years after the hearing (in granted cases) or to no longer have an order (in denied cases).

Our restraining order request sample does contain some limited information on the characteristics of the request. In addition to knowing the outcome of the RO request at the hearing (granted vs. denied), our data also includes information on the presence of lawyers for either party, whether the case involved a minor child, and the type of restraining order requested.[22] We also know if either party requested a continuance, whether the petitioner confirmed service with the court before the hearing, if the petitioner notified the court that service failed, or if the court noted that the respondent had not yet been notified.

The summary statistics for our sample of restraining order requests are shown in Table 2. The breakdown by different case types, those denied at the hearing, and those granted at the hearing, suggests that the composition of cases differs meaningfully across categories. For example, cases in which the order is denied at the 21-day hearing are less likely to involve legal representation or confirmed service and are less likely to involve continuances. By contrast, cases that proceed to hearing and result in a grant show much higher rates of confirmed service and somewhat higher involvement of minor children. Lawyer presence overall is rare across all subsamples, but slightly more common when petitions are granted.

3.2 LAPD Crime Data

Our data on crimes committed at the block level address comes from the LAPD dataset, "Crime Data from 2020 to Present," which provides detailed records of reported crimes within the City of Los Angeles.[23]

[21] Typically, all address information, correspondence address, and banned address were missing from the request if there was no banned address. Court personnel have told us that, because restraining order requests are public documents, legal aid and judicial officers recommend that in orders where the respondent does not know the petitioner's address, the petitioner should not provide it and give the court a way to communicate with the petitioner that does not reveal their address to the respondent, typically a phone number. The exception is when the petitioner has a lawyer, whose address is provided for correspondence purposes, which is rare in our dataset.

[22] While lawyers are uncommon in restraining order cases, their presence appears to impact case outcomes substantially. Our results are robust to excluding cases with lawyers.

[23] Crime Data from 2020 to Present:
`https://data.lacity.org/Public-Safety/Crime-Data-from-2020-to-Present/2nrs-mtv8`

It captures incidents recorded by the Los Angeles Police Department (LAPD) regardless of whether they resulted in an arrest. In addition to allowing us to examine a complete measure of deterrence, this database is more suitable for our analysis than the arrest data used in previous studies. The arrest data provided by the LAPD only contains the final policing actions, thereby obscuring civilian reports. In the City of Los Angeles, the crime dataset also serves as the basis for the official domestic violence calls database, which has been defined and provided by the LAPD since 2020, allowing us to measure "domestic violence" classifications, a feature not available in the arrest data.[24] The LAPD data includes crime type, location (mapped to neighborhood and reporting district), date and time, and limited information on victim demographics. As of 2024, the database had over 1.3 million entries and was updated daily by the LAPD.

The LAPD rounds the address data to the nearest block.[25] Figure 1 shows a hypothetical example of the block-address match. Black stars indicate the address that may be listed in the restraining order request, while blue squares denote the size of a block. A hashed line represents a hypothetical crime committed on the block in the year following the grant or denial of the restraining order. In our sample, the observation at the far left of West 35th PL, which has a hypothetical restraining order indicated by the blue star, would be included in our data with the dependent variable equal to one, indicating that a crime had been committed on the block in the year following the 21-day hearing. By contrast, the observation directly to the right has a restraining order petition but no crime. The boxes on West 36th Street, directly south, would not be included in our analysis because no address in our sample had a restraining order on that block.[26]

One concern is that the block observation level potentially introduces a mismatch between the address in the restraining order and the location of crimes. As discussed below, we include crime incidents from January 2020 through December 2024, and our analysis is limited to residential crimes classified in the data as domestic violence or assaults, totaling 134,000 incidents. There are 51,800 unique block locations based on coordinates in the LAPD data. The average number of incidents per location is 2.59 (134,000 / 51,800 = 2.59). Thus, over the four years, the average number of incidents per block location per year is 2.59 / 4 = 0.65. In addition, with around 5,000 restraining order petitions per year, the number of petitions per block per year is 0.1.

To put the mismatch problem in perspective, consider the likelihood that a crime on the block is randomly distributed across addresses on the block. The number of unique addresses on a block in Los Angeles varies widely, depending on neighborhood density, zoning, and the method of address assignment.[27] Assuming a rough upper bound of an average of 60 addresses per block and a lower bound of 30, based on only our rough count of addresses in our sample blocks, we can estimate the probability that a crime incident recorded at

[24] Domestic Violence Calls from 2020 to Present:
`https://data.lacity.org/Public-Safety/Domestic-Violence-Calls-from-2020-to-Present/qq59-f26t/about_data`

[25] For instance, 1021 Childs Way is rounded to 1000 Childs Way.

[26] In our sample, we have no blocks with overlapping restraining orders.

[27] For example, certain blocks have larger lot sizes or undeveloped parcels, while others might have many apartment buildings or multi-unit dwellings in which residents have individual addresses.

the block level occurs randomly at the specific address associated with the restraining order is between 1.67 and 3.33%.

The low probability of a crime incident, a restraining order at the block level, and the high probability of mismatch if crimes were randomly distributed suggest that there is a considerable amount of "noise" generated by the mismatch between the block and the actual address. In other words, detecting a genuine relationship between restraining orders and subsequent crime becomes much harder if most incidents on the block do not occur at the house under the restraining order. Consequently, the underlying relationship is likely robust if we find a consistent and statistically significant effect of restraining orders on crime outcomes despite this noise. If mismatches were systematically biasing the data, we would expect the signal of any real effect to be drowned out. The fact that the effect still emerges, despite the relatively small probability of an exact match, indicates the effect is unlikely to be spurious.[28]

Importantly for our study, the LAPD data contains information on whether the investigating officer labeled the incident as domestic violence. Broadly, a domestic violence crime is identified based on the relationship between the victim and the suspect. These crimes typically involve violent or threatening behavior within a domestic relationship, such as between spouses, intimate partners, cohabitants, or family members. Domestic violence crimes can encompass any offense and are not limited to those offenses typically thought of as domestic violence, e.g., assault, battery, stalking, threats, etc. The LAPD also cautions that this classification is at the discretion of the investigating officer, so underreporting is a concern.

We confront several tradeoffs in determining which crimes to use as our dependent variable. For our preferred specification, we focus on crimes explicitly labeled as domestic violence by the LAPD. One concern with this approach is that, in addition to potential underreporting, police officers may be more likely to classify a crime as domestic violence if the victim has received a ROAH. Since all approved orders are reported to the police, who have access to their provisions, classification could be influenced, but only if the order was actually granted. To the extent that this occurs, it would bias our estimates toward zero, attenuating the true effect of restraining orders on crime. To address this concern, we estimate the model with several alternative measures of crime.

We refer to these alternatives as the "assault continuum." The continuum of criminal assaults in the context of domestic violence reflects a progression of severity, beginning with criminal threats (verbal or written threats to harm), escalating to simple assault (minor physical harm or attempted harm), aggravated assault (serious injuries or use of a weapon), and battery (intentional and harmful physical contact), and culminating in murder. We also include child abuse and rape within this definition. Although this measure captures the vast majority of serious domestic violence incidents, it is narrower than our preferred measure because it

[28] Using a leniency design IV also mitigates concerns about mismatch. Because judge assignment is random, the causal effect of a restraining order on crime is identified independently of any potential mismatch in the location-level data. Thus, any mismatch is likely to introduce noise rather than systematic bias.

excludes offenses such as property destruction and certain types of harassment that are broadly recognized as forms of domestic violence. The key advantage of the assault continuum is that reporting of these serious, clearly defined crimes is unlikely to be influenced by the existence of a restraining order, thereby mitigating concerns about treatment-induced measurement bias.

To further narrow the measure to crimes covered by the restraining order, we focus only on residential crimes. The data also contains a standardized description of the location of each crime. The most common premises descriptions are "single-family dwelling" (17%), "multi-unit dwelling" (12%), and "street" (25%). To increase the likelihood of capturing crimes occurring at the location specified in the restraining order request, we include only premises descriptions that broadly correspond to residential locations in the sample.[29]

One final problem is that the willingness to report a crime may be correlated with order status. For example, if the petitioner is denied the order, they may be less likely to report a crime. Although this would tend to bias us toward finding no effect on the probability or number of crimes, it is worth considering how our two approaches to addressing the endogeneity of granting orders with respect to crime would also address these issues.

We estimate a leniency design, where judges' tendencies to grant restraining orders serve as an instrumental variable. Given that judges lack the ability to follow up with the implementation of the order, judge leniency is plausibly exogenous to the crime-reporting behavior and other confounders. Importantly, the random assignment of cases to judges ensures that there is no correlation between judge leniency and unobserved factors that influence reporting behavior. As with all leniency designs, this approach assumes that the instrument (judge leniency) affects crime only through granting restraining orders and not through other pathways (e.g., judges contacting police to ask them to more carefully monitor the petitioner's address after the order is granted).

4 Research Design

To address the endogeneity of an order being granted and the probability of a crime, we utilize a judge leniency/examiner design that leverages the differences in the probability that a restraining order request is granted across judges. Given that judges have broad discretion in deciding whether to grant a restraining order request, it is perhaps not surprising to see a wide variation across judges, and the likelihood that a judge grants the request appears to have a judge-specific component.

A distinctive feature of our setting is that the Stanley Mosk Courthouse, which handles domestic violence cases, adjudicates cases from across Los Angeles County, whereas our outcome variable is restricted to

[29]Specifically, we include only crimes with a premises description of Single-Family Dwelling, Multi-Unit Dwelling (Apartment, Duplex, etc.), Garage/Carport, Driveway, Other Residence, Yard (Residential/Business), Residential Porch, Condominium/Townhouse, Laundromat, Mobile Home/Trailer/Construction Trailer/RV/Motorhome, Transitional Housing/Halfway House, Patio, Single Residence Occupancy (SRO) Locations, Project/Tenement/Public Housing, Balcony, or Apartment/Condo Common Laundry Room. 92% of the residential crimes occur in single-family (48%), multi-unit (34%), garages (5%), or driveways (5%).

incidents within the City of Los Angeles. Over our sample period (2020–2023), we observe approximately 30,000 cases county-wide (29,831 for our baseline residualized leniency measure). Restricting the judge leniency measure to the geographic boundaries of the outcome variable would not only discard roughly one-third of the available observations but would also likely introduce selection bias. We therefore construct the instrument using the full set of approximately 30,000 county-wide observations.

4.1 Judge Leniency Design

This section outlines our research design for estimating the causal effect of granting a restraining order at the 21-day hearing on the probability that a residential crime classified by the LAPD as residential domestic violence occurs at the block address listed in the restraining order petition within the 365-day window following the hearing. The outcome is denoted $Pr(\text{crime})_i$, and the unit of observation is the petition i.

$$\Pr(\text{crime})_i = \alpha_0 + \alpha_1 \,\text{grant}_i + \mathbf{X}_i\gamma + \delta_{ct} + \delta_y + \delta_m + \delta_d + \varepsilon_i$$

Here, grant_i indicates whether the restraining order petition for individual i, which was not withdrawn before the 21-day hearing, was granted, and α_1 is our parameter of interest. The key challenge is that an OLS estimate of α_1 may confound two effects: a selection effect (petitions more likely to involve future violence are also more likely to be granted) and a deterrence effect (grants reduce subsequent crime).

The vector $\mathbf{X}_i$ that includes whether the docket lists evidence of proof of service, whether minor children are involved, whether the petitioner had an attorney, whether the respondent had an attorney, and race probabilities inferred from names. In addition, we included fixed effects for year, δ_y, month of the year, δ_m, day of the week, δ_d, and petition type δ_{ct}. ε_i are the standard errors clustered by filing date, since, as explained above, this is the level at which randomization occurs (Abadie et al., 2023; Chyn et al., 2025). We also provide Anderson–Rubin 95% confidence intervals. The treatment in this sample is moving from having a restraining order in the form of the TRO to either continuing or not having a restraining order after the hearing.

The second-stage estimating equation corresponds to a standard instrumental variables setup, in which we use judge leniency as an instrument for whether a restraining order request is issued. Formally, the first stage is a linear probability model:

$$\text{grant}_i = \alpha_0 + \alpha_1 Z_{ij} + \mathbf{X'}_i\gamma + \varepsilon_i$$

where Z_i is our residualized Jackknife Instrumental Variables Estimator (rJIVE) of judge leniency. $\mathbf{X'}_i$ contains the battery of fixed effects, described above.

Following Angrist et al. (1999) and Chyn et al. (2025), we construct a leave-one-out measure of examiner leniency that accounts for structured calendar variation. Because restraining order requests may vary systematically by year, month, or day of the week, a simple leave-one-out measure of judge leniency could

be biased by such scheduling patterns (Agan et al. 2023). We therefore define examiner leniency Z_{ij} as the residualized, leave-one-out mean probability that judge j grants the restraining order. To account for case-mix and temporal scheduling patterns, we first residualize the grant decision on fixed effects for petition type, filing year, filing month, and day of week:

$$\text{grant}^*_{ij} = \text{grant}_{ij} - \mathbf{X'}_i\gamma = Z_{ij} + \varepsilon_i,$$

where $\mathbf{X'}_i$ includes a battery of fixed effects, case type, year, month, and day-of-week fixed effects. The residual grant decision, grant^*_{ij}, includes our measure of judge leniency Z_{ij}, as well as idiosyncratic case-level variation ε_i.

The measure is constructed across all restraining order requests, not only those linked to address data, and we exclude judges with fewer than 5 cases to ensure stability. Judge j's leniency for case i is defined as the leave-one-out average of these residualized grant outcomes among their other cases:

$$Z_{ij} = \left(\frac{1}{n_j - 1}\right)\left(\sum_{i'=0}^{n_j} \text{grant}^*_{i'j} - \text{grant}^*_{ij}\right),$$

where n_j is the number of cases seen by judge j. By removing the specific case in question, i, we ensure that variation in Z_{ij} reflects relative judge leniency, net of case-specific influence.

Ideally, following Dobbie et al. (2018), after residualizing the grant rate using case type and time fixed effects, we would calculate the JIVE estimator using only the current institutional year, consistent with the institutional setting in which practitioners rotate in and out on an annual basis. However, this approach leads to a substantial loss of sample size (Table 1), making the measure noisier. We therefore begin by including the previous year in addition to the current year. Incorporating all prior years yields robust results.[30]

After removing judges who hear fewer than five cases per year, the sample includes 90 judges. The median number of petitions heard is 51.5, while the average is 191.59, consistent with a small number of judges in the standalone restraining order courts hearing the majority of petitions. After residualizing for time effects, the judge's leniency measures vary from -.1359 at the first percentile and .3188 for the ninety-ninth percentile of judge leniency. With a mean grant rate in the sample of .248, moving from the first to the ninety-ninth percentile of judge leniency changes the grant rate by .1654 percentage points or 66.74 %.

Frandsen et al. (2023) identify four key conditions for causal interpretation of judge- or examiner-based instrumental variable (IV) designs:

1. First-stage relevance: the instrument (judge assignment) must be correlated with the treatment; in our context, a judge's overall granting propensity must predict whether a given petition is granted.

[30]Pooling all years forward, however, produces a noisier measure, likely due to the frequent rotations characteristic of the family court context, see above.

2. As-good-as-random assignment: cases or petitions must be quasi-randomly assigned to judges, as occurs in the LASC.

3. Exclusion restriction: judge assignment affects the outcome only through its impact on the treatment (granting the restraining order), and not through any other channel of judicial behavior that could directly influence outcomes such as crime rates.

4. Monotonicity: the effect of judge assignment on treatment is weakly monotonic—i.e., if one judge is more lenient than another, every case that would be granted by the stricter judge would also be granted by the more lenient one. Violations occur if judges differ in which types of cases they are lenient toward.

To evaluate the first stage, Table 4 presents the first stage coefficient and F-statistics using four versions of the judge leniency instrument across measures. Chyn et al. (2025) argue that the first-stage coefficient in an examiner design is critical for assessing the instrument's strength, i.e., how much variation in treatment is explained by the instrument. A higher value means a stronger influence of the instrument on treatment assignment, which improves IV identification (see Angrist and Kolesár (2024)). Because JIVE type estimators, of the kind that are typically employed in examiner/judge designs and that we use here, are more sensitive to weak instruments than 2SLS Angrist and Kolesár (2024) recommend a first-stage coefficient of at least 0.64 under a one-sided test of weak instrument concerns. As shown in Table 4, our first-stage coefficients are .65 and .69 for the residualized JIVE estimator, depending on whether we include covariates and depending on the sample period used to construct the estimator. This exceeds the threshold for reliability under a one-sided test, indicating that our JIVE estimator is unlikely to suffer from weak instrument bias. The first stage F statistics are 146.587, 144.845, and 134.046, all far in excess of the value of 10 used to identify weak instruments.[31]

The assignment process outline above results in diverse approaches to handling cases. Given that the evidentiary standard for issuing a restraining order under the Domestic Violence Prevention Act (DVPA) is "reasonable proof of a past act or acts of abuse," interpreted as a preponderance of the evidence standard, we might expect to find a wide variation in grant rates, which is what we observe in Figure 2.

In our case, the exclusion restriction would require that judges who are more likely to grant restraining orders do not also intervene in ways that increase the probability of incidents or arrests. Discussions with judicial officers at the Los Angeles Superior Court (LASC) suggest there is no direct contact between judges and the LAPD beyond the issuance of the order, making it unclear how judges could influence police behavior. A potential concern is that the existence of a protective order might increase the likelihood of an arrest; however,

[31] When using a leave-one-out constructed instrument based on examiner or judge fixed effects, the usual first-stage F-statistic from regressing the endogenous variable on the single constructed instrument will be mechanically inflated relative to the "true" first-stage F from the many-dummy specification. As (Hull, 2017) shows, the inflation factor is approximately equal to the number of examiner groups under homoskedasticity, so that weak instruments may be masked.

such an effect would bias our analysis against finding any impact of protective orders, thereby mitigating concerns about overestimating their effects. Moreover, our estimates are based on reported crimes rather than arrests, further mitigating this concern.

Moreover, our discussions with judicial officers also suggest that judges' interactions with petitioners typically begin and end with the hearings on restraining orders. Even if they wanted to intervene in some way to reduce the likelihood of a crime against the petitioner, judges' ability to influence the likelihood that the petitioner is the victim of a crime is essentially nonexistent. Judges lack mechanisms for follow-up after issuing an order. Once an order is issued, it is entered into a police databank, but judges themselves have no ability to contact the LAPD or monitor enforcement. In fact, during our sample period, judges in civil restraining order courts were unable to access the criminal records of respondents. Moreover, firearm surrender compliance is limited to a signed affidavit from the respondent stating that the weapon has been sold or stored. Unless the petitioner explicitly denies compliance or the affidavit is missing, there is no police involvement.

Frandsen et al. (2023)'s final criterion is the monotonicity requirement, which states that if Judge A is more likely to grant a restraining order than Judge B, every petition granted by Judge B would also have been granted by Judge A. In our context, monotonicity would be violated if, for instance, one judge were lenient only in certain types of petitions (e.g., domestic violence without minor children) while another were lenient across types.

Frandsen et al. (2023) develop a joint test of the two core assumptions underlying the validity of leniency-based instrumental variables: monotonicity (that more "lenient" judges always treat weak cases that stricter judges would not) and exclusion (that judges affect outcomes only through the treatment). A statistically significant result indicates that at least one of the two assumptions, monotonicity or exclusion, is violated. Table 3 reports the Chi-squared statistics from this test for our sample and residualized leniency measures across four subsamples. In Panel A, we reject the null hypothesis that both assumptions hold for the full sample when using residualized JIVE with and without controls. In other words, at least one of the assumptions are not supported. Similarly, when focusing on a specific case type, civil harassment orders (Panel B), we again reject the null that both monotonicity and exclusion hold. For all DVROs (Panel C) and the subsample of DVRO petitions with minor children (Panel D) are we unable to reject the null that both monotonicity and exclusion hold. When we examine petitions without children (Panel E) we again reject the null, suggesting that at least one identifying assumption is violated.

Given the implausibility of judges affecting crime at an address through channels other than the issuance of a restraining order, we suspect the violations arise from monotonicity. This is plausible in our context, as judges with very low grant rates often handle only a small number of petitions, increasing the likelihood of idiosyncratic treatment reversals.

In our context, where the treatment, granting a restraining order, is binary and judge assignment is random, a failure of strict monotonicity suggests that treatment effects may not be point-identified for compliers, and thus marginal treatment effects are not recoverable. However, our estimates can still be interpreted as average treatment effects over a broader subset of petitioners under weaker conditions such as average monotonicity and average exclusion. This implies that even if the FLL assumptions fail, the IV estimates remain meaningful but reflect average rather than marginal effects.

Figure 2 shows the residualized grant rates for the full sample (top panel) and for DVROs with children (bottom panel). The bar graph shows the distribution of the judge leniency measure, and the line graph shows the correlation between the judge leniency measure (on the x-axis) and the probability of granting a petition request (on the y-axis), which is the graphical representation of the first stage of the instruments. The y-axis on the left lays out the density of the distribution, and the y-axis on the right shows the probability that a petition is granted. We use the residualized grant rate and leniency measure to determine the as-if random conditions across claim types and years.

For the monotonicity condition to hold, we should see the probability of granting a petition increase throughout the distribution of judges. We see this in Figure 2. Specifically, it shows an increasing probability that a petition will be granted (on the y-axis) as we move up the distribution of judges (on the x-axis). The estimate of grant probability is imprecise at the distribution's tails due to the lack of observations. The gray lines indicate the 95% confidence interval, and the widening range of the gray lines as we approach a grant rate of 1 means that the correlation between the judge's leniency and the probability of granting a specific petition is unreliable at the highest grant rates.

5 Results

Table 4 presents the regression results when we pool all order types, beginning with the OLS specification and progressing to our two instrumental variable strategies. In column (1) of Panel A, we estimate the model without controlling for endogeneity and without including control variables on the probability of a police-identified domestic violence crime. In column (2), we estimate the model including our control variables. In both columns (1) and (2), we find a small, negative, and significant coefficient. The impact of granting a restraining order is small, reducing the probability of a crime by 1.31 percentage points in column (1) and by 1.30 percentage points in column (2). To put this in perspective, consider the relative risk reduction calculation in the bottom row of panel A. In the row labeled 'implied percent effect,' we show the ratio of the impact of granting the restraining order divided by the average probability that a crime will be committed at the block-level address over the next year, as measured by the relevant crime measure. The elasticity measures the intervention's proportional effect relative to the baseline probability of the outcome.

To further characterize the selection process, we estimate the share of compliers following Abadie (2003) (see

also Bhuller et al. (2020) and Agan et al. (2023)). In an instrumental variables framework, Abadie (2003) shows that the extremes of the instrument distribution identify always-takers (those who would receive treatment even under the strictest judge, $\overline{z}$) and never-takers (those who would not receive treatment even under the most lenient judge, $\underline{z}$), with the difference between these two groups identifying the share of compliers, $\overline{z} - \underline{z}$. Using this method, we estimate that 21.6% of the sample are compliers.

Following Agan et al. (2023), Column (3) reports complier-weighted OLS estimates, reweighting the sample to approximate the distribution of observed characteristics among compliers. The estimate is closer to zero than the conventional OLS estimate, although the two are not statistically distinguishable. Thus, the observable characteristics used in the reweighting provide limited evidence that the complier population differs sharply from the overall analysis sample.

Before presenting the IV estimates, we report the reduced-form relationship between judicial leniency and subsequent crime. The reduced-form estimates are negative across all three specifications: assignment to a judge with a higher residualized propensity to grant restraining orders is associated with a lower probability of a police-identified domestic-violence crime during the following year. The estimated coefficients range from -0.035 to -0.050, although they are imprecisely estimated in the pooled sample. These estimates capture the effect of assignment to a more lenient judge on subsequent crime, prior to rescaling by the effect of judicial leniency on the probability that the order is granted. The corresponding IV estimates in Columns (4)–(6) divide the reduced-form coefficients by the first-stage coefficients reported at the bottom of the table.

In Columns (4)-(6), we rescale these reduced-form effects by the corresponding first stages using our residualized JIVE instrument. The estimates are larger in magnitude than the OLS estimates, but remain imprecise in the pooled sample. In Column (4), granting a restraining order is estimated to reduce the probability of a police-identified domestic-violence crime at the block-level address by 7.44 percentage points. Relative to the sample mean, this corresponds to an implied reduction of approximately 45 percent. Including the full set of controls produces a similar estimate of 6.69 percentage points. Constructing the residualized JIVE instrument using all prior years yields a somewhat smaller estimate and does not improve precision.

Panel B uses the broader assault-continuum outcome. The OLS estimates are small, while the implied reduced-form coefficients range from approximately -0.049 to -0.060. The corresponding IV estimates imply reductions of 7.60 to 8.88 percentage points, but the confidence intervals include zero. Panel C focuses on aggravated assault. Here, the OLS and complier-weighted OLS estimates are close to zero and slightly positive, whereas the implied reduced forms and IV estimates are negative. The IV estimates range from -3.43 to -4.05 percentage points and are statistically insignificant. Panel D shows a similar pattern for simple assaults: the implied reduced forms are negative, ranging from approximately -0.023 to -0.037, and the IV estimates range from -3.55 to -5.45 percentage points, but remain imprecisely estimated.

The pooled estimates provide suggestive evidence of negative effects but do not allow us to reject a zero av-

erage effect across all restraining-order petitions. The difference between the OLS and IV point estimates is nevertheless consistent with substantial selection into treatment. As emphasized by Bjerk (2025), differences between OLS and IV estimates may reflect either selection into treatment or treatment-effect heterogeneity. In this setting, judges may be more likely to grant orders in cases involving a greater underlying risk of future violence, which would bias the OLS estimate toward zero or even toward a positive association. The similarity between conventional and complier-weighted OLS estimates suggests that differences in the observed characteristics used for reweighting do not explain much of the OLS–IV gap. It does not, however, rule out differences on unobservables or heterogeneous treatment effects. We therefore interpret the comparison as consistent with adverse selection into granting, rather than as a definitive decomposition of the OLS-IV difference.

A second reason the pooled estimates may obscure meaningful effects is that granting a "restraining order" does not represent a homogeneous legal intervention across petition types. Although violations of the different orders carry similar criminal penalties, domestic-violence restraining orders involving minor children can also determine custody, visitation, possession of the residence, child support, and spousal support. These accompanying orders may sharply reduce the contact that would otherwise be legally or financially difficult to avoid. Civil-harassment orders and domestic-violence restraining orders without minor children generally do not provide the same bundle of family-law remedies. If these additional provisions are important to deterrence, pooling all petitions will average across substantively different treatments. In the next section, we estimate the model separately by restraining-order type.

5.1 Breakdown by Type of Order Requested

Table 5 disaggregates the results by restraining-order type, using police-identified residential domestic-violence crime as the outcome. As shown in Table 2, the unconditional probability of the outcome is broadly similar across the different petition categories. The table reports OLS estimates, implied reduced-form relationships between judicial leniency and subsequent crime, and the corresponding rJIVE estimates.

Panel A examines civil-harassment orders, which are generally filed against individuals with whom the petitioner does not have a close domestic relationship. Both the OLS and rJIVE point estimates are small and positive,and the Anderson–Rubin confidence intervals include zero. The implied reduced-form coefficients are also positive, ranging from 0.011 to 0.027. Thus, we find no evidence that assignment to a more lenient judge, or the granting of a civil-harassment order itself, reduces subsequent police-identified domestic-violence crime at the petitioner's block.

Panel B includes all domestic-violence restraining-order petitions. The OLS estimates are negative but modest in magnitude. The implied reduced-form coefficients are consistently negative, ranging from approximately -0.105 to -0.120, indicating that assignment to a judge with a higher propensity to grant DVROs is associated with a lower probability of subsequent crime. After rescaling by the first stage, the rJIVE esti-

mates imply reductions of approximately 14.1 to 15.2 percentage points. These estimates are substantially larger in magnitude than the OLS estimates, but the Anderson–Rubin confidence intervals include zero.

Panel C focuses on DVRO petitions involving minor children. Here, the OLS estimates are close to zero and slightly positive, whereas both the implied reduced forms and the rJIVE estimates are large and negative. The implied reduced-form coefficients are remarkably stable across specifications, ranging from -0.259 to -0.265. The corresponding rJIVE estimates imply that granting a DVRO reduces the probability of a police-identified residential domestic-violence crime by approximately 33.7 to 36.3 percentage points. The estimates are statistically significant using conventional clustered standard errors, and the Anderson–Rubin confidence intervals exclude zero.[32]

Panel D examines DVRO petitions without minor children. The OLS estimates are negative, while the implied reduced-form coefficients are much closer to zero, ranging from approximately -0.007 to -0.027. The corresponding rJIVE estimates range from -0.9 to -3.4 percentage points, and all Anderson–Rubin confidence intervals include zero. Thus, the evidence of a deterrent effect is concentrated in DVRO petitions involving minor children.

The pattern is consistent with the possibility that DVROs involving minor children constitute a different legal treatment bundle. In addition to restricting contact, these orders may determine custody, visitation, possession of the residence, child support, and spousal support. Those accompanying provisions can reduce contact that would otherwise be legally or financially difficult to avoid.

Table 6 examines whether the estimated effects for DVRO petitions involving minor children extend to alternative measures of crime. Panel A repeats the police-defined domestic violence estimates from Table 5 for comparison. As discussed above, the estimated impact is large, negative, and statistically significant across all three rJIVE specifications.

Panel B broadens the outcome to include the full assault continuum. The point estimates remain uniformly negative, and the implied reduced-form coefficients are also consistently negative. Although the estimated effects are somewhat smaller than those for police-defined domestic violence, the confidence intervals include zero. Thus, the broader assault measure is consistent in sign with the primary outcome, but is estimated less precisely.

Panel C focuses exclusively on aggravated assaults. Once again, the point estimates are negative across all specifications and considerably larger in magnitude than those for the broader assault measure. While none of the rJIVE estimates are individually significant, the estimates are directionally consistent with the primary findings and suggest that any reduction in violent crime may be concentrated among the more serious offenses.

[32]Because these point estimates exceed the unconditional mean of the binary outcome, expressing them as percentage changes relative to the overall sample mean produces values greater than 100 percent. We therefore interpret the estimates in percentage-point terms rather than literally as percentage reductions. Economically, the estimates imply that granting a DVRO nearly eliminates the observed block-level crime outcome for the marginal petitions affected by judicial assignment.

Panel D examines simple assaults, threats, and child abuse after excluding aggravated assaults. Here, the estimated effects are smaller and fluctuate around zero, with confidence intervals that comfortably include both economically meaningful deterrence effects and no effect. We therefore find little evidence that the reductions observed for police-defined domestic violence are driven by changes in less serious assault offenses. The evidence indicates that the clearest and most robust effects are found for police-defined domestic violence among petitions involving minor children. The remaining crime measures generally display point estimates that are similar in direction but substantially less precise. Because aggravated assaults are relatively rare events, the available sample provides limited power to detect effects of plausible magnitude. Accordingly, we view these estimates primarily as providing supportive evidence regarding the pattern of treatment effects rather than as independent tests of the intervention's effectiveness.

5.2 Robustness Check: Excluding the COVID-19 Remote Hearing Period

One potential concern with our study period (January 1, 2020, through December 31, 2023) is that it begins during the COVID-19 pandemic, when restraining-order hearings in Los Angeles were temporarily conducted remotely. Although the Los Angeles Superior Court classified restraining-order proceedings as essential and continued hearing petitions by video or telephone, remote hearings could have altered judicial decision-making, petitioner participation, or compliance with court orders.[33]

To examine this possibility, Table 7 re-estimates our preferred specification after excluding petitions filed between March 15 and June 5, 2020, the period during which the Stanley Mosk restraining-order courts operated remotely. The estimated effects are very similar to those obtained from the full sample. The implied reduced-form relationship between judicial leniency and subsequent crime remains negative, and the corresponding rJIVE estimates continue to imply that granting a domestic violence restraining order reduces the probability of a police-defined residential domestic-violence incident by approximately 33 to 36 percentage points. The Anderson–Rubin confidence intervals continue to exclude zero. Thus, our principal findings do not appear to be driven by the unusual court procedures adopted during the initial months of the COVID-19 pandemic.

5.3 Robustness Check: Alternative Instrument Construction

Table 8 examines whether our findings are robust to alternative methods of constructing the judge-leniency instrument. Our preferred specification uses the residualized jackknife IV (rJIVE), which removes the mechanical correlation between each case and the judge-specific grant rate used to construct the instrument. This section asks whether our substantive conclusions depend on that particular implementation.

Columns (1) and (2) estimate the model using conventional two-stage least squares with judge fixed effects as instruments, following Agan et al. (2023), who employ a similar first-stage specification. This

[33]Legal aid services also remained available remotely, with limited in-person assistance.

approach exploits the same underlying source of variation as the rJIVE estimator but does not incorporate the leave-one-out jackknife correction. Consistent with our preferred specification, the estimated effects remain negative. The point estimates imply reductions of approximately 17 to 19 percentage points in the probability of a police-defined domestic violence incident. Under conventional clustered standard errors, the estimates are statistically significant at the 10 percent level, while a cluster bootstrap using 894 filing-date clusters yields significance at the 1 percent level.

Columns (3)-(6) implement two alternative jackknife estimators: Influence-function JIVE (IJIVE) and Unbiased JIVE (UJIVE). IJIVE, suggested by Chyn et al. (2025) and Agan et al. (2023), combines leave-one-out estimation with an influence-function correction to reduce finite-sample bias. UJIVE, introduced by Angrist and Kolesár (2024), removes the remaining finite-sample bias analytically and has attractive theoretical properties when many instruments are available. Across both estimators, the estimated effects remain remarkably similar to those obtained with our preferred residualized JIVE specification. Point estimates range from approximately -0.30 to -0.31, indicating that the substantive conclusions are not sensitive to the particular jackknife estimator employed. Although the confidence intervals are somewhat wider than those for the preferred rJIVE specification, the point estimates remain highly stable across estimators.

Finally, in Columns (7) and (8), we estimate the model using the cluster-jackknife instrumental variables estimator (CJIVE) proposed by Frandsen et al. (2025). Whereas the leave-one-out estimators remove the mechanical correlation between an observation and its own judge-specific leniency measure, CJIVE extends this logic to settings in which observations within a cluster may also contribute to finite-sample bias. In our application, the natural clustering unit is the calendar day, since restraining-order hearings are conducted in daily court sessions, and our primary inference clusters standard errors at the filing-date level. We therefore compute the judge-specific leniency measure after excluding all observations from the same filing-date cluster. Because judges who are not assigned to the standalone restraining-order court on a given day may nevertheless have been available during that court session, we construct the estimator using the full roster of judges. In practice, however, judges outside the dedicated restraining-order court rarely hear restraining-order petitions. The CJIVE estimates are similar in sign, magnitude, and statistical significance to the preferred residualized JIVE estimates. The alternative instrument constructions produce remarkably similar estimates despite relying on different bias corrections and identifying assumptions, suggesting that the principal findings are not driven by the particular implementation of the judge-leniency instrument.

5.4 Robustness Check: Saturated Time Residualization

Table 9 examines whether our results are sensitive to the manner in which we residualize the judge-leniency measure with respect to calendar time. Our baseline specification removes additive time effects before constructing the residualized JIVE instrument. One concern is that if judicial assignment or grant rates vary along more complex combinations of calendar effects, an additive specification may not fully absorb

those sources of variation.

To address this possibility, Table 9 reports estimates from a saturated residualized JIVE specification in which the time fixed effects are fully interacted during the residualization step. The estimated effects remain highly similar to those obtained from the baseline specification, ranging from -0.337 to -0.363. The consistency of the estimates across alternative temporal windows used to construct the judge-leniency measure indicates that our results are robust to alternative specifications of the time adjustment.

5.5 Robustness Check: Statistical Power

One possible explanation for the heterogeneous treatment effects across restraining-order types is statistical power. Because the estimated effects for civil-harassment orders and domestic violence restraining orders without minor children are imprecisely estimated, it is possible that meaningful deterrent effects exist, but cannot be detected in those subsamples.

To assess this possibility, Table 10 reports minimum detectable effects (MDEs) for each subsample using a back-of-the-envelope calculation adapted to our judge-IV setting. The calculation combines the cluster-robust reduced-form standard error with the estimated first-stage coefficient to obtain the smallest local average treatment effect detectable with 80 percent power at the 5 percent significance level. Appendix A provides the derivation.

The estimated MDEs are remarkably similar across the four restraining-order categories. The reduced-form MDEs range from approximately 0.20 to 0.32, while the corresponding IV MDEs range from 0.32 to 0.45. Thus, the statistically insignificant estimates for civil harassment petitions and DVRO petitions without minor children cannot be attributed simply to substantially lower statistical power. Instead, the results suggest that the large deterrent effects are concentrated among DVRO petitions involving minor children.

This pattern is consistent with the institutional differences across petition types. Unlike the other restraining orders, DVROs involving minor children frequently operate as bundled legal interventions, combining no-contact provisions with custody, visitation, child support, spousal support, and exclusive possession orders. These accompanying family-law remedies may substantially reduce continued interaction between the parties, providing a plausible mechanism for the heterogeneous treatment effects documented above.

6 Conclusion

Using administrative records from the Los Angeles Superior Court linked to crime reports, this paper estimates the causal effect of granting a civil restraining order on subsequent residential domestic violence. Because judges differ systematically in their propensity to grant restraining orders, ordinary least squares comparisons confound the effect of the order itself with differences in the underlying risk faced by petitioners. We address this problem using a judge-leniency instrument based on residualized leave-one-out jackknife

estimation.

Pooling all restraining-order petitions, we find little evidence that granting an order reduces subsequent crime. Disaggregating the analysis, however, reveals substantial treatment heterogeneity. The deterrent effect is concentrated entirely among domestic violence restraining order petitions involving minor children. For civil harassment petitions, domestic violence petitions without minor children, and the pooled sample, we find no statistically significant deterrent effect despite strong first stages and comparable statistical power.

The institutional setting suggests a natural explanation for this heterogeneity. Although all restraining orders prohibit contact and carry criminal penalties for violations, a granted domestic violence restraining order involving minor children frequently bundles additional legal remedies that substantially alter the parties' ongoing relationship. These remedies may include temporary sole custody, visitation restrictions (often supervised), exclusive possession of the residence, child support, and spousal support. Collectively, these provisions reduce ongoing interaction between the petitioner and respondent and lower the cost of self-protection compared with cases in which the order is denied. By contrast, petitioners in the other categories are not legally constrained from using alternative means of reducing contact even without a restraining order, making the marginal effect of judicial granting substantially smaller.

The robustness analyses reinforce this interpretation. The principal results are unchanged when excluding the COVID remote-hearing period, using alternative jackknife estimators, employing saturated time residualization, or comparing statistical power across petition types, suggesting that the observed heterogeneity is unlikely to be an artifact of the estimation procedure or differences in sample size.

More broadly, our findings suggest that civil restraining orders should not be viewed as a single legal intervention. The effectiveness of a restraining order may depend not only on the legal prohibition against future contact, but also on the complementary bundle of legal rights and obligations that accompanies the order. For domestic violence petitions involving minor children, granting a restraining order changes not only the legal sanctions for future abuse but also the family's custody, housing, and financial arrangements. Evaluating the effectiveness of restraining-order policy, therefore, requires recognizing that these accompanying family-law remedies may be at least as important as the restraining order itself.

Our instrumental variables estimates identify a Local Average Treatment Effect (LATE): the effect of granting a restraining order for petitions whose disposition is influenced by judicial assignment. Under the conditions described by Heckman and Vytlacil (2001), this quantity may also provide useful guidance for evaluating institutional reforms that shift judicial granting practices. Such reforms are directly relevant in practice. Courts regularly provide judicial training, revise bench guides, and reconsider the evidentiary standards governing restraining-order hearings. Our results, therefore, speak not only to the consequences of an individual judge's decision but also to the potential effects of systematic changes in how courts evaluate restraining-order petitions.

Our findings also underscore the difficult policy tradeoff facing family courts. Granting a domestic violence restraining order may substantially reduce the risk of future violence, but it can also alter custody, visitation, and other aspects of the parent-child relationship. Conversely, denying an order preserves those relationships but may leave petitioners and their children exposed to continued abuse. This paper cannot determine the socially optimal threshold for granting restraining orders. Instead, it demonstrates that the tradeoff is economically meaningful and that the bundled legal remedies accompanying domestic violence restraining orders involving minor children appear to play a central role in their effectiveness.

References

Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of econometrics 113*(2), 231–263.

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics 138*(1), 1–35.

Agan, A., J. L. Doleac, and A. Harvey (2023). Misdemeanor prosecution. *The Quarterly Journal of Economics 138*(3), 1453–1505.

Aizer, A. and P. Dal Bo (2009). Love, hate and murder: Commitment devices in violent relationships. *Journal of public Economics 93*(3-4), 412–428.

Amaral, S., G. B. Dahl, V. Endl-Geyer, T. Hener, and H. Rainer (2023). Deterrence or backlash? arrests and the dynamics of domestic violence. Technical report, National Bureau of Economic Research.

Anderson, T. W. and H. Rubin (1949). Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of mathematical statistics 20*(1), 46–63.

Angrist, J. and M. Kolesár (2024). One instrument to rule them all: The bias and coverage of just-id iv. *Journal of Econometrics 240*(2), 105398.

Angrist, J. D., G. W. Imbens, and A. B. Krueger (1999). Jackknife instrumental variables estimation. *Journal of Applied Econometrics 14*(1), 57–67.

Berk, R. A., A. Campbell, R. Klap, and B. Western (1992). A bayesian analysis of the colorado springs spouse abuse experiment. *J. Crim. L. & Criminology 83*, 170.

Bhuller, M., G. B. Dahl, K. V. Løken, and M. Mogstad (2020). Incarceration, recidivism, and employment. *Journal of political economy 128*(4), 1269–1324.

Bjerk, D. (2025). The coefficient difference check for evaluating iv results.

Black, M. C., K. C. Basile, M. J. Breiding, S. G. Smith, M. L. Walters, M. T. Merrick, J. Chen, and M. R. Stevens (2011). The national intimate partner and sexual violence survey: 2010 summary report.

Carlson, M. J., S. D. Harris, and G. W. Holden (1999). Protective orders and domestic violence: Risk factors for re-abuse. *Journal of Family Violence 14*(2), 205–206.

Chin, Y.-M. and S. Cunningham (2019). Revisiting the effect of warrantless domestic violence arrest laws on intimate partner homicides. *Journal of Public Economics 179*, 104072.

Chyn, E., B. Frandsen, and E. Leslie (2025). Examiner and judge designs in economics: a practitioner's guide. *Journal of Economic Literature 63*(2), 401–439.

Djimeu, E. W. and D.-G. Houndolo (2016). Power calculation for causal inference in social science: sample size and minimum detectable effect determination. *Journal of Development Effectiveness 8*(4), 508–527.

Dobbie, W., J. Goldin, and C. S. Yang (2018). The effects of pretrial detention on conviction, future crime, and employment: Evidence from randomly assigned judges. *American Economic Review 108*(2), 201–240.

Duflo, E., R. Glennerster, and M. Kremer (2007). Using randomization in development economics research: A toolkit. *Handbook of development economics 4*, 3895–3962.

Dunford, F. W., D. Huizinga, and D. S. Elliott (1990). The role of arrest in domestic assault: The omaha police experiment. *Criminology 28*(2), 183–206.

Frandsen, B. (2020). Testjfe: Stata module to perform test for instrument validity in the judge fixed effects design. Statistical Software Components S458871, Boston College Department of Economics, revised 31 Jan 2023.

Frandsen, B., L. Lefgren, and E. Leslie (2023). Judging judge fixed effects. *American Economic Review 113*(1), 253–277.

Frandsen, B., E. Leslie, and S. McIntyre (2025). Cluster jackknife instrumental variables estimation. *Review of Economics and Statistics*, 1–19.

Graef, L., S. G. Mayson, A. Ouss, and M. T. Stevenson (2023). Systemic failure to appear in court. *U. Pa. L. Rev. 172*, 1.

Harper, S. B., A. R. Gover, and T. N. Richards (2021). Protection orders: Shielding intimate partner violence victims from harm. In *Handbook of Interpersonal Violence and Abuse Across the Lifespan: A project of the National Partnership to End Interpersonal Violence Across the Lifespan (NPEIV)*, pp. 3021–3041. Springer.

Heckman, J. J. and E. Vytlacil (2001). Policy-relevant treatment effects. *American Economic Review 91*(2), 107–111.

Hirschel, J. D. and I. W. Hutchison iII (1992). Female spouse abuse and the police response: The charlotte, north carolina experiment. *J. Crim. L. & Criminology 83*, 73.

Hull, P. (2017). Examiner designs and first-stage f statistics: A caution. Unpublished Working Paper.

Iyengar, R. (2009). Does the certainty of arrest reduce domestic violence? evidence from mandatory and recommended arrest laws. *Journal of public Economics 93*(1-2), 85–98.

Judicial Council of California (2023, October). *Judges Guide to Domestic Violence Restraining Orders*. Judicial Council of California. With special acknowledgment to Ms. Sarah Jacobvitz, Esq., Hon. Susan M. Gill, Hon. Mark A. Juhas.

Klein, A. (1996). Re-abuse in a population of court-restrained male batterers: Why restraining orders don't work. In E. Buzawa and C. Buzawa (Eds.), *Do Arrests and Restraining Orders Work?*, pp. 192–210. Thousand Oaks, CA: Sage Publications.

Ko, C. N. (2001). Civil restraining orders for domestic violence: The unresolved question of efficacy. *S. Cal. Interdisc. LJ 11*, 361.

Los Angeles County Department of Public Health, Injury and Violence Prevention Program (2013, February). Quick facts about intimate partner violence (ipv). Los Angeles County Department of Public Health, Injury and Violence Prevention Program.

Pate, A. M. and E. E. Hamilton (1992). Formal and informal deterrents to domestic violence: The dade county spouse assault experiment. *American Sociological Review*, 691–697.

Sherman, L. W. and R. A. Berk (2017). The specific deterrent effects of arrest for domestic assault. In *Quantitative Methods in Criminology*, pp. 3–14. Routledge.

Sherman, L. W., J. D. Schmidt, D. P. Rogan, and D. A. Smith (1992). The variable effects of arrest on criminal careers: The milwaukee domestic violence experiment. *J. Crim. L. & Criminology 83*, 137.

Tjaden, P. G. (2000). *Full report of the prevalence, incidence, and consequences of violence against women: Findings from the National Violence Against Women Survey*. US Department of Justice, Office of Justice Programs.

Table 1: Case Type among Analysis Sample Across Years (Los Angeles, 2020–2023)

Case type	2020	2021	2022	2023	Total
Civil Harassment Prevention	2,046	2,149	2,372	2,850	9,417
DV Prevention w/ Minor Child	651	745	801	949	3,145
DV Prevention w/o Minor Child	1,158	1,329	1,711	1,961	6,159
Elder/Dependent Adult Abuse Prevention	257	308	307	409	1,281
Elder/Dependent Adult Restraining Order	0	0	0	1	1
Gun Violence Restraining Order	3	4	4	21	32
Workplace Violence Prevention	7	13	36	38	94
Total	4,122	4,548	5,231	6,229	20,130

Table 2: Descriptive Statistics of Restraining Order Cases (Los Angeles, 2020–2024)

	Full Analysis Sample			ROAH Denied			ROAH Granted		
	Average	Min	Max	Average	Min	Max	Average	Min	Max
Any lawyer presence	0.061	0	1	0.054	0	1	0.080	0	1
Lawyers for both parties	0.001	0	1	0.001	0	1	0.002	0	1
Lawyer for petitioner	0.034	0	1	0.027	0	1	0.056	0	1
Lawyer for respondent	0.036	0	1	0.037	0	1	0.034	0	1
Lack of Prosecution	0.514	0	1	0.632	0	1	0.151	0	1
Continuance	0.340	0	1	0.338	0	1	0.355	0	1
Service failed once	0.027	0	1	0.025	0	1	0.034	0	1
Service confirmed	0.507	0	1	0.415	0	1	0.805	0	1
Minor child involved	0.154	0	1	0.144	0	1	0.196	0	1
Domestic Violence (Panel A)	0.166	0	1	0.168	0	1	0.157	0	1
Panel A Subset - Civil Harassment	0.155	0	1	0.156	0	1	0.149	0	1
Panel A Subset - DVRO with child	0.183	0	1	0.181	0	1	0.186	0	1
Panel A Subset - DVRO w/o child	0.181	0	1	0.191	0	1	0.158	0	1
Assault Continuum (Panel B)	0.225	0	1	0.226	0	1	0.220	0	1
Aggravated Assault (Panel C)	0.059	0	1	0.059	0	1	0.060	0	1
Simple Assault (Panel D)	0.165	0	1	0.167	0	1	0.161	0	1
Observations	20384			15328			4803		

Notes: This table summarizes data from 20,384 restraining order petitions from the city of Los Angeles, including 253 cases missing judge information.

Table 3: Test Statistics for Joint Null Hypotheses of Exclusion and Monotonicity (Frandsen et al., 2023)

	No controls	All controls
Panel A. Entire Observations		
χ^2 Statistic	148.80	154.70
fit-based p-value	0.003	0.001
Combined p-value	0.006	0.002
Panel B. Civil Harassment Order		
χ^2 Statistic	126.64	142.01
fit-based p-value	0.012	0.001
Combined p-value	0.023	0.002
Panel C. DVRO (All)		
χ^2 Statistic	109.51	121.96
fit-based p-value	0.242	0.067
Combined p-value	0.484	0.134
Panel D. DVRO with Minor Children		
χ^2 Statistic	74.56	101.40
fit-based p-value	0.894	0.214
Combined p-value	1.000	0.428
Panel E. DVRO without Minor Children		
χ^2 Statistic	145.35	144.44
fit-based p-value	0.001	0.001
Combined p-value	0.001	0.001

Notes: The analysis controls for the occurrence of a continuance, the outcome of service (fulfilled or failed), legal representation of the petitioner, respondent, or both parties, and race probabilities inferred from names, in addition to fixed effects for year, month, and day of the week. The test is conducted using the `testjfe` program in Stata, as described in Frandsen (2020).

Table 4: OLS Regression and Instrumental Variable Strategy Crimes One Year Post Filing

	(1) OLS	(2) OLS	(3) C-OLS	(4) rJIVE	(5) rJIVE	(6) rJIVE
Panel A. Police-defined domestic violence						
RO Grant	-0.0131*	-0.0130^{+}	-0.00624	-0.0744	-0.0669	-0.0535
	(0.00630)	(0.00661)	(0.00768)	(0.0725)	(0.0737)	(0.0771)
AR 95% CI				[-.220,.068]	[-.215,.078]	[-.208,.098]
Implied Percent Effect	-0.079	-0.078		-0.449	-0.404	-0.323
Implied Reduced Form				-0.050	-0.045	-0.035
Panel B. All Assault Continuum Incidents						
RO Grant	-0.00740	-0.00571	0.000247	-0.0888	-0.0853	-0.0760
	(0.00738)	(0.00779)	(0.00913)	(0.0826)	(0.0838)	(0.0871)
AR 95% CI				[-.255,.073]	[-.254,.079]	[-.252,.095]
Implied Percent Effect	-0.033	-0.025		-0.395	-0.380	-0.338
Implied Reduced Form				-0.060	-0.057	-0.049
Panel C. Assault Continuum – Aggravated						
RO Grant	0.000807	0.00144	0.000647	-0.0343	-0.0352	-0.0405
	(0.00407)	(0.00432)	(0.00510)	(0.0463)	(0.0469)	(0.0507)
AR 95% CI				[-.126,.058]	[-.129,.058]	[-.142,.060]
Implied Percent Effect	0.014	0.024		-0.579	-0.593	-0.683
Implied Reduced Form				-0.023	-0.023	-0.026
Panel D. Assault Continuum – Simple						
RO Grant	-0.00821	-0.00715	-0.000401	-0.0545	-0.0501	-0.0355
	(0.00648)	(0.00692)	(0.00820)	(0.0722)	(0.0730)	(0.0750)
AR 95% CI				[-.200,.087]	[-.197,.093]	[-.186,.112]
Implied Percent Effect	-0.050	-0.043		-0.330	-0.303	-0.215
Implied Reduced Form				-0.037	-0.033	-0.023
Observations	20130	20079	18082	20129	20078	20080
Control Variables	No	Yes	Yes	No	Yes	Yes
JIVE Years				Current&Last	Current&Last	Current&All Prior
First stage F-stat				146.587	144.845	134.046
First stage coefficient				0.672	0.665	0.650
Complier Weights	No	No	Yes			

Notes: This table reports the main regression results by specification, excluding cases dismissed before the restraining-order hearing. Columns (1) and (2) report OLS specifications. Column (3) reports complier-weighted OLS estimates, splitting the sample into 10 mutually exclusive groups based on deciles of the predicted probability that an order is granted, estimated using all available covariates, following Agan et al. (2023). Columns (4) and (5) report instrumental-variable estimates using residualized JIVE measures constructed from the current and previous year. Column (6) uses a residualized JIVE measure constructed from the current and all prior years. The implied reduced-form coefficients are calculated as the product of the reported rJIVE estimate and the corresponding first-stage coefficient. Because these reduced-form coefficients are reconstructed rather than estimated directly, standard errors are not reported. Each panel uses a different crime outcome. Panel A focuses on residential crimes classified by police as domestic violence. Panel B includes the assault continuum. Panel C focuses on aggravated assaults within Panel B. Panel D includes simple assaults, threats, and child abuse from Panel B, excluding the aggravated assaults used in Panel C. The analysis controls for the occurrence of a continuance, the outcome of service (fulfilled or failed), legal representation of the petitioner, respondent, or both parties, and race probabilities inferred from names, in addition to fixed effects for year, month, and day of the week. Standard errors, clustered by filing date, are reported in parentheses. Anderson–Rubin 95% confidence intervals are shown in square brackets.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $0.05 \leq p < 0.10$.

Table 5: Results by Restraining Order Type

	(1) OLS	(2) OLS	(3) rJIVE	(4) rJIVE	(5) rJIVE
Panel A: Civil Harassment Order					
RO Grant	-0.00428	-0.00705	0.0212	0.0358	0.0530
	(0.0103)	(0.0110)	(0.138)	(0.136)	(0.138)
AR 95% CI			[-.265,.299]	[-.247,.307]	[-.232,.333]
Implied Percent Effect	-0.028	-0.046	0.137	0.231	0.343
Implied Reduced Form			0.011	0.018	0.027
Observations	9417	9406	9416	9405	9406
First stage F-stat			44.985	49.864	48.861
First stage Coefficient			0.531	0.511	0.515
Panel B: DVRO					
RO Grant	-0.0194*	-0.0167+	-0.141	-0.152	-0.139
	(0.00853)	(0.00906)	(0.0907)	(0.0928)	(0.0980)
AR 95% CI			[-.325,.038]	[-.340,.031]	[-.338,.055]
Implied Percent Effect	-0.107	-0.092	-0.777	-0.835	-0.767
Implied Reduced Form			-0.114	-0.120	-0.105
Observations	9305	9297	9305	9297	9298
First stage F-stat			97.335	98.029	86.324
First stage Coefficient			0.810	0.787	0.755
Panel C: DVRO with minor children					
RO Grant	0.00375	0.00386	-0.340*	-0.337*	-0.363*
	(0.0151)	(0.0158)	(0.167)	(0.166)	(0.184)
AR 95% CI			[-.743,-.039]	[-.734,-.038]	[-.824,-.037]
Implied Percent Effect	0.020	0.021	-1.859	-1.843	-1.987
Implied Reduced Form			-0.265	-0.262	-0.259
Observations	3146	3142	3146	3142	3142
First stage F-stat			33.992	36.108	28.913
First stage Coefficient			0.779	0.776	0.714
Panel D: DVRO without minor children					
RO Grant	-0.0319**	-0.0282*	-0.0250	-0.0341	-0.00926
	(0.0107)	(0.0116)	(0.107)	(0.112)	(0.117)
AR 95% CI			[-.231,.203]	[-.250,.205]	[-.219,.259]
Implied Percent Effect	-0.176	-0.156	-0.138	-0.188	-0.051
Implied Reduced Form			-0.021	-0.027	-0.007
Observations	6159	6155	6159	6155	6156
First stage F-stat			63.760	63.452	58.418
First stage Coefficient			0.825	0.791	0.774
Control Variables	No	Yes	No	Yes	Yes
JIVE Years			Current&Last	Current&Last	Current&All Prior

Notes: This table replicates the main regression results from Panel A of Table 4, police-defined domestic violence calls, disaggregating the results by order type and excluding cases that are not Domestic Violence Restraining Orders (DVROs) or civil restraining orders. The analysis controls for the occurrence of a continuance, the outcome of service (fulfilled or failed), legal representation of the petitioner, respondent, or both parties, and race probabilities inferred from names, in addition to fixed effects for year, month, and day of the week. Standard errors, clustered by filing date, are reported in parentheses. Anderson–Rubin 95% confidence intervals are shown in square brackets.
*** p<0.001, ** p<0.01, * p<0.05, + 0.10 p<0.1.

Table 6: Focusing on Domestic Violence Restraining Order with Minor Child

	(1) OLS	(2) OLS	(3) rJIVE	(4) rJIVE	(5) rJIVE
Panel A: Police-defined domestic violence					
RO Grant	0.00375	0.00386	-0.340*	-0.337*	-0.363*
	(0.0151)	(0.0158)	(0.167)	(0.166)	(0.184)
AR 95% CI			[-.743,-.039]	[-.734,-.038]	[-.824,-.037]
Implied Percent Effect	0.020	0.021	-1.859	-1.843	-1.987
Implied Reduced Form			-0.265	-0.262	-0.259
Panel B: All Assault Continuum Incidents					
RO Grant	0.00579	0.0104	-0.279	-0.285	-0.287
	(0.0166)	(0.0174)	(0.176)	(0.178)	(0.196)
AR 95% CI			[-.687,.051]	[-.700,.043]	[-.758,.075]
Implied Percent Effect	0.024	0.043	-1.154	-1.183	-1.188
Implied Reduced Form			-0.217	-0.221	-0.205
Panel C: Assault Continuum – Aggravated					
RO Grant	-0.00120	0.000205	-0.121	-0.124	-0.148
	(0.00938)	(0.00985)	(0.0860)	(0.0853)	(0.0971)
AR 95% CI			[-.315,.044]	[-.317,.038]	[-.376,.036]
Implied Percent Effect	-0.020	0.003	-2.036	-2.084	-2.488
Implied Reduced Form			-0.094	-0.096	-0.106
Panel D: Assault Continuum – Simple					
RO Grant	0.00699	0.0102	-0.158	-0.162	-0.139
	(0.0152)	(0.0161)	(0.152)	(0.153)	(0.167)
AR 95% CI			[-.498,.137]	[-.504,.131]	[-.520,.184]
Implied Percent Effect	0.038	0.056	-0.866	-0.889	-0.763
Implied Reduced Form			-0.123	-0.126	-0.099
Observations	3146	3142	3146	3142	3142
Control Variables	No	Yes	No	Yes	Yes
JIVE Years			Current&Last	Current&Last	Current&All Prior

Notes: This table replicates the regression results from Table 4, focusing exclusively on domestic violence restraining-order petitions involving minor children and excluding civil-harassment petitions and domestic violence restraining-order petitions without minor children. Panel A repeats the preferred specification from Panel C of Table 5. The analysis controls for the occurrence of a continuance, the outcome of service (fulfilled or failed), legal representation of the petitioner, respondent, or both parties, and race probabilities inferred from names, in addition to fixed effects for year, month, and day of the week. The implied reduced-form coefficients are calculated as the product of the reported rJIVE estimate and the corresponding first-stage coefficient reported in Table 5. Because these reduced-form coefficients are reconstructed rather than estimated directly, standard errors are not reported. Standard errors, clustered by filing date, are reported in parentheses. Anderson–Rubin 95% confidence intervals are shown in square brackets.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $0.05 \leq p < 0.10$.

Table 7: Robustness to COVID-19 Court Closures (Domestic Violence Restraining Order with Minor Child)

	(1) OLS	(2) rJIVE	(3) rJIVE
RO Grant	0.00399	-0.333^{+}	-0.355^{+}
	(0.0161)	(0.179)	(0.196)
AR 95% CI		[-.764,-.011]	[-.852,-.007]
Implied Percent Effect	0.022	-1.813	-1.936
Implied Reduced Form		-0.258	-0.253
Observations	3045	3045	3045
Control Variables	Yes	Yes	Yes
JIVE Years	Current&Last	Current&All Prior	

Notes: This table reproduces the preferred specification from Table 6, Panel A, after excluding petitions filed between March 15, 2020, when restraining-order hearings moved to virtual proceedings because of the COVID-19 pandemic, and June 5, 2020, when the dedicated restraining-order courts resumed in-person hearings. Unlike the baseline specification, this robustness check retains the full set of judges rather than excluding peripheral judges who heard fewer than five restraining-order hearings during the previous year. All specifications control for the occurrence of a continuance, the outcome of service (fulfilled or failed), legal representation of the petitioner, respondent, or both parties, race probabilities inferred from names, and fixed effects for year, month, and day of the week. The implied reduced-form coefficients are calculated as the product of the reported rJIVE estimate and the corresponding first-stage coefficient from the preferred specification. Because these reduced-form coefficients are reconstructed rather than estimated directly, standard errors are not reported. Standard errors, clustered by filing date, are reported in parentheses. Anderson–Rubin 95% confidence intervals are shown in square brackets.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $0.05 \leq p < 0.10$.

Table 8: Alternative Instrument Measures (Domestic Violence Restraining Order with Minor Child)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Dummy	Dummy	IJIVE	IJIVE	UJIVE	UJIVE	CJIVE	CJIVE
RO Grant	-0.175**	-0.194**	-0.296+	-0.285	-0.301+	-0.313+	-0.379*	-0.415+
	(0.062)	(0.066)	(0.161)	(0.183)	(0.170)	(0.184)	(0.191)	(0.216)
AR 95% CI	[-.396,.026]	[-.439,.024]	[-.682,-.004]	[-.739,.048]	[-.705,.008]	[-.768,.019]	[-.871,-.042]	[-1.014,-.042]
Observations	3142	3142	3142	3142	3142	3142	3142	3142
Average Elasticity	-0.959	-1.060	-1.621	-1.557	-1.646	-1.714	-2.076	-2.270
Control Variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
JIVE Years	Current&Last	Current&All Prior	Current&Last	Current&All Prior	Current&Last	Current&All Prior	Current&Last	Current&All Prior

Notes: This table replicates the regression results from Table 6, Panel A, using alternative methods to construct the instrument. Columns (1)–(2) use judge indicators as instruments; Columns (3)–(4) use IJIVE; Columns (5)–(6) use UJIVE; and Columns (7)–(8) use CJIVE. All specifications include controls for the occurrence of a continuance, the outcome of service (fulfilled or failed), legal representation of the petitioner, respondent, or both parties, and race probabilities inferred from names, in addition to fixed effects for year, month, and day of the week. Standard errors, clustered by filing date, are reported in parentheses. Anderson–Rubin 95% confidence intervals are shown in square brackets. For Columns (1)–(2), standard errors are computed using a cluster bootstrap in which 894 daily filing-date clusters are resampled with replacement to account for the fitted-instrument setting with many judge indicators as instruments.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $0.05 \leq p < 0.10$.

Table 9: Saturated Residualized JIVE (Domestic Violence Restraining Order with Minor Child)

	(1) rJIVE	(2) rJIVE	(3) rJIVE	(4) rJIVE
RO Grant	-0.312^{+}	-0.331^{+}	-0.322^{+}	-0.297
	(0.175)	(0.195)	(0.179)	(0.205)
AR 95% CI	[-.737,.004]	[-.829,.016]	[-.757,.009]	[-.818,.086]
Average Elasticity	-1.707	-1.811	-1.761	-1.623
Observations	3142	3142	3105	3107
Control Variables	Yes	Yes	Yes	Yes
JIVE Years	Current&Last	Current&All Prior	365 Days < Today	All < Today

Notes: This table replicates the regression results from Table 6, focusing on Panel A, but employs fully interacted time fixed effects by year, month, and day of the week when residualizing the leniency measure. The analysis controls for the occurrence of a continuance, the outcome of service (fulfilled or failed), legal representation of the petitioner, respondent, or both parties, and race probabilities inferred from names, in addition to fixed effects for year, month, and day of the week. Standard errors, clustered by filing date, are reported in parentheses. Anderson–Rubin 95% confidence intervals are shown in square brackets.
*** p<0.001, ** p<0.01, * p<0.05, + 0.10 p<0.1.

Table 10: Power and Minimum Detectable Effects for Judge-IV Design

Subsample	N	SE_{ITT}	π (First Stage)	MDE_{ITT}	MDE_{IV}
Civil Harassment	9,406	0.072	0.512	0.201	0.392
DVRO with Child	3,142	0.118	0.724	0.329	0.455
DVRO without Child	6,155	0.091	0.779	0.255	0.327

Notes: ITT stands for Intention-to-Treat effect, the reduced-form coefficient of outcome on instrument. SE_{ITT} is the cluster-robust standard error from the reduced form (incident on instrument) using the same controls and fixed effects as the IV. π is the first-stage coefficient (grant on instrument). MDE_{ITT} is computed as $(z_{\text{power}} + z_{1-\alpha/2}) \times SE_{ITT}$ with power = 0.80 and $\alpha = 0.05$; MDE_{IV} equals $MDE_{ITT}/|\pi|$.

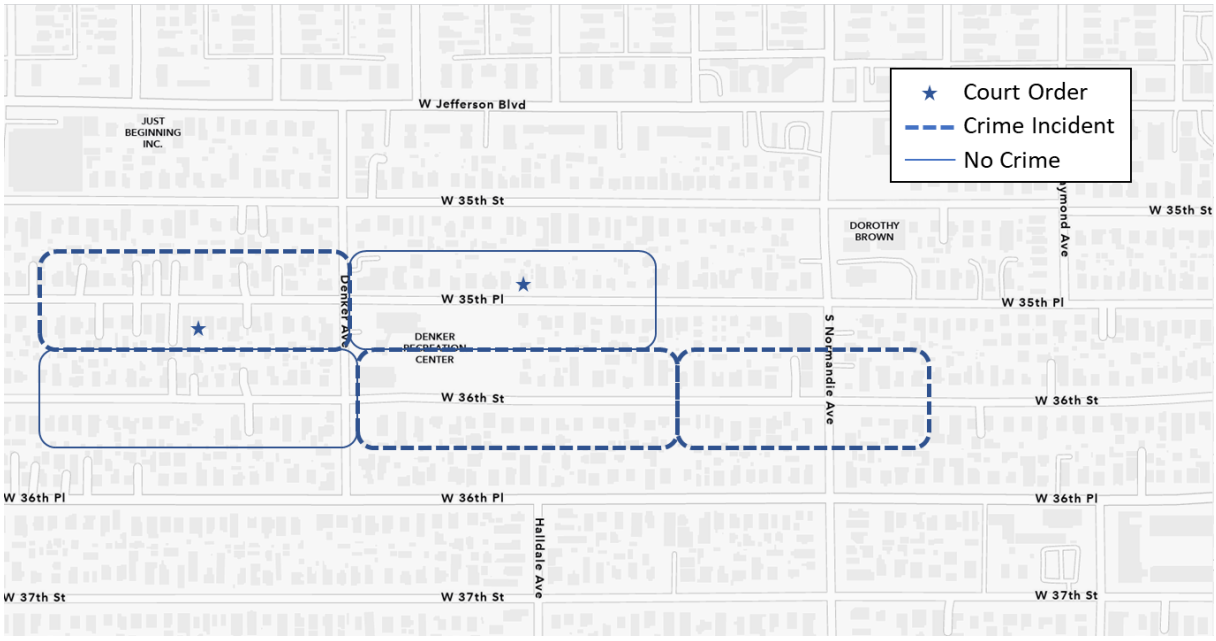


Figure 1: Hypothetical Examples of 100-block Address Match

Note: This figure illustrates the matching process for addresses aggregated to the 100-block level. Black stars are hypothetical addresses similar to those found in restraining order requests, blue boxes show 100 blocks, and dashed lines indicate a hypothetical crime on that street.

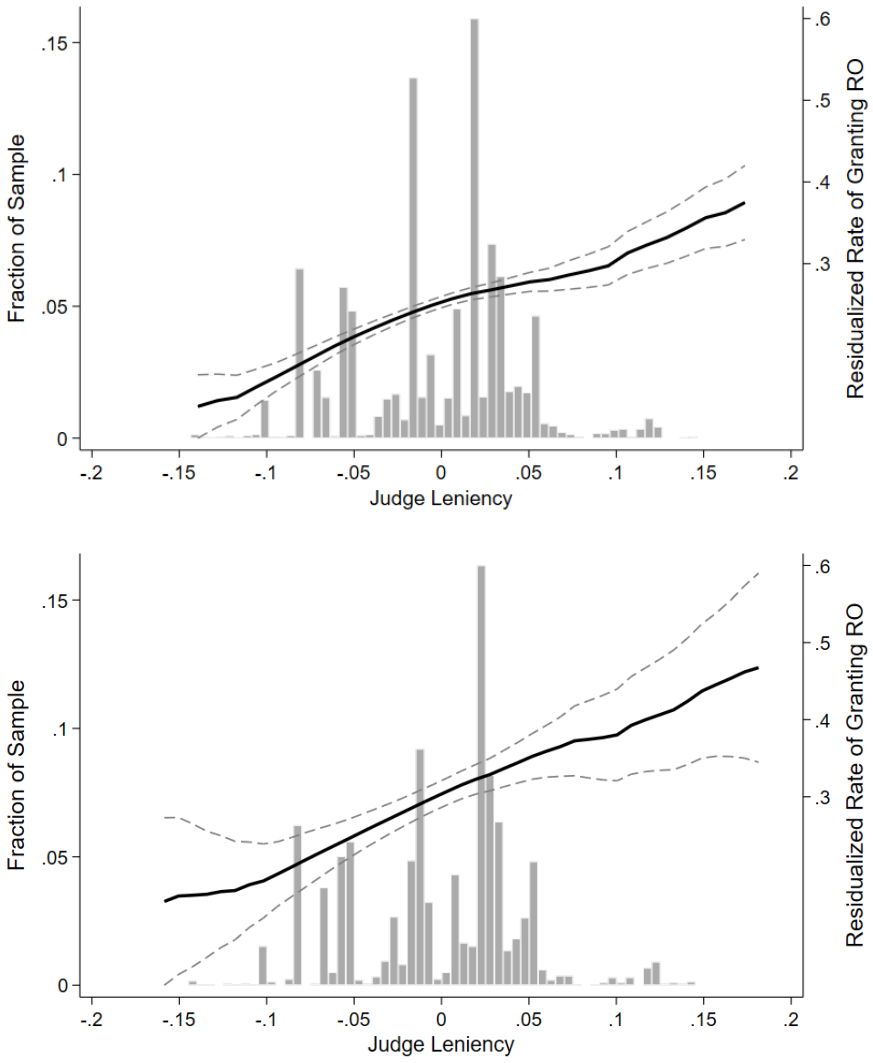


Figure 2: Distribution of Residualized Leniency Measure and First Stage

Note: This figure presents the distribution of judge leniency measures and their relationship with petition grant rates. The top panel shows the residualized grant rates for the full sample and the lower panel for DVROs involving children. In each panel, the bar graph shows the distribution of the judge leniency measure, while the line graph plots the probability of granting a petition request (y-axis) against the judge leniency measure (x-axis). The left y-axis displays the density of the leniency measure, and the right y-axis indicates the probability that a petition is granted. Gray lines represent the 95% confidence intervals. The residualized leniency measure is obtained by regressing the grant rate on year-, month-, and day-of-week fixed effects and petition characteristics and then averaging the residuals for each judge.

A Appendix: Minimum Detectable Effects for Judge-IV Design

This appendix describes the calculation used to obtain the minimum detectable effects (MDEs) reported in Table 10. Our objective is not to provide an exact finite-sample power calculation for the instrumental variables estimator, but rather to compare the amount of statistical power available across the different restraining-order subsamples.

Following Duflo et al. (2007) and Djimeu and Houndolo (2016), the minimum detectable effect is the smallest true treatment effect that can be detected with a specified significance level and statistical power. For a simple randomized experiment, the conventional expression is

$$MDE = (t_{1-\beta} + t_{1-\alpha/2})\sqrt{\frac{1}{P(1-P)}\frac{\sigma^2}{N}}, \tag{1}$$

where α denotes the probability of a Type I error, β denotes the probability of a Type II error, P is the treatment share, σ^2 is the outcome variance, and N is the sample size.

Our setting differs from this canonical framework in two respects. First, treatment is not randomly assigned but instrumented by judge leniency. Second, inference relies on cluster-robust standard errors rather than the homoskedastic variance formula above. Consequently, rather than using the analytical variance from the experimental design literature, we base the calculation directly on the estimated cluster-robust reduced-form standard error.

Let SE_{RF} denote the cluster-robust standard error from the reduced-form regression of the outcome on the judge-leniency instrument. Assuming asymptotic normality, the minimum detectable reduced-form effect is

$$MDE_{RF} = (z_{1-\beta} + z_{1-\alpha/2})SE_{RF}. \tag{2}$$

Throughout the paper, we assume 80 percent statistical power and a 5 percent two-sided significance level, implying

$$z_{1-\beta} = 0.84, \qquad z_{1-\alpha/2} = 1.96,$$

so that

$$z_{1-\beta} + z_{1-\alpha/2} = 2.80.$$

Finally, because the instrumental variables estimator rescales the reduced-form coefficient by the first-stage coefficient, the implied minimum detectable local average treatment effect is

$$MDE_{IV} = \frac{MDE_{RF}}{|\pi|}, \tag{3}$$

where π denotes the estimated first-stage coefficient.

This calculation is intended as a diagnostic measure rather than an exact finite-sample power analysis. Because every subsample uses the same estimator, clustering scheme, and inference procedure, the resulting MDEs provide a useful comparison of the amount of statistical power available across restraining-order types. The similarity of the estimated MDEs indicates that the heterogeneous treatment effects reported in the main text are unlikely to be explained solely by differences in statistical power.